



Pin-Yu Chen · Sijia Liu

Introduction to Foundation Models

 Springer

Introduction to Foundation Models

Pin-Yu Chen • Sijia Liu

Introduction to Foundation Models

 Springer

Pin-Yu Chen
IBM Research
Yorktown Heights, NY, USA

Sijia Liu
College of Engineering
Michigan State University
Sunnyvale, CA, USA

ISBN 978-3-031-76769-2 ISBN 978-3-031-76770-8 (eBook)
<https://doi.org/10.1007/978-3-031-76770-8>

© The Editor(s) (if applicable) and The Author(s), under exclusive license to Springer Nature Switzerland AG 2025

This work is subject to copyright. All rights are solely and exclusively licensed by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

If disposing of this product, please recycle the paper.

*To a better future and technology where AGI
means Artificial Good Intelligence*

Preface

Back in June 2022, we (Pin-Yu Chen, Sijia Liu, and Sayak Paul) sensed that foundational models and generative AI were about to bring unprecedented and revolutionary changes to our technology and society, and we decided to propose a tutorial at NeurIPS to educate researchers and practitioners about this emerging AI research and technology and its implications for robustness and safety. Two months later, our tutorial, “[Foundational Robustness of Foundation Models](#),” was accepted and to be presented on December 5, 2022. At that time, the world was still recovering from the Covid crisis, and NeurIPS 2022 was the first NeurIPS conference to allow physical attendance after Covid. To accommodate the hybrid conference mode, we were asked to record our tutorial by October 2022. In our materials, we correctly predicted that the proliferation of foundation models would bring increased risks and challenges in robustness, security, and safety, later known as part of the broad topic of AI safety, governance, and risk assessment.

About a week before the official release of our tutorial and the start date of NeurIPS 2022, ChatGPT was launched by OpenAI on November 30, 2022. At the conference, we heard only a few whispers and genuine inquiries about “Have you heard of ChatGPT?” Who would have thought that in August 2024, the number of ChatGPT users worldwide would be more than 180 million?

With such a coincidence in mind, we felt compelled to expand our tutorial into a book, summarizing the basics of foundation models and generative AI, and emphasizing the robustness and safety of these frontier models and AI technology.

This book consists of three parts. Part I (Chaps. 1 to 4) provides the fundamentals of foundation models, Part II (Chaps. 5 to 11) includes advanced topics in foundation modes, and Part III (Chaps. 12 to 18) presents safety and trust in foundation models. The mathematical notations will be defined and explained in each chapter.

We provide an overview of each part as follows:

- In Part I, Chap. 1 provides an overview of foundation models and generative AI. Chapter 2 presents the technical background of neural networks. Chapter 3 delves into the learning and generalization of transformers. Chapter 4 formalizes in-context learning with transformers.

- In Part II, Chap. 5 introduces automated visual prompting techniques. Chapter 6 introduces prompting LLMs with privacy. Chapter 7 elucidates memory-efficient fine-tuning methods. Chapter 8 shows how LLMs can be reprogrammed for time-series machine learning tasks. Chapter 9 shows how LLMs can be reused for speech tasks. Chapter 10 introduces how synthetic datasets can be used to benchmark foundation models. Chapter 11 elucidates machine unlearning for foundation models.
- In Part III, Chap. 12 provides a comprehensive evaluation of the trustworthiness of LLMs. Chapter 13 introduces jailbreak attacks and defenses for LLMs. Chapter 14 presents safety risks when fine-tuning LLMs. Chapter 15 introduces watermarking techniques for LLMs. Chapter 16 presents robust detection of AI-generated text. Chapter 17 elucidates backdoor risks in diffusion models. Chapter 18 presents red-teaming methods for diffusion models.

The authors would especially like to thank Paul Drougas at Springer for his encouragement and patience in allowing us to publish a book in such a rapidly changing and growing field. Some parts of the book are based on publications co-authored by the authors of this book. We would like to express our gratitude to our research collaborators. Finally, we would like to thank our families, colleagues, and friends, for their support.

Yorktown Heights, NY, USA
Sunnyvale, CA, USA
September 2024

Pin-Yu Chen
Sijia Liu

Contents

Part I Fundamentals of Foundation Models

1	Foundation Models and Generative AI	3
1.1	What Are Foundation Models and Generative AI?	3
1.2	Foundations Models at Scale	5
1.2.1	Neural Scaling Laws	5
1.2.2	Emerging Abilities	6
1.3	Lifecycle of Foundation Models	7
1.3.1	Data Preparation	8
1.3.2	Model Training	8
1.3.3	Model Deployment	9
1.4	Canonical Examples of Foundation Models	10
1.5	Our Perspective on Foundation Models	11
2	Neural Networks	13
2.1	Introduction	13
2.2	Basic Neural Network Modules	14
2.3	Transformers	16
2.3.1	Token Embedding and Position Encoding	16
2.3.2	Attention	17
2.3.3	Transformer Types	19
2.4	Large Language Models	20
2.4.1	Next-Token Prediction	20
2.4.2	Decoding Strategies	21
2.4.3	Alignment Strategies	22
2.4.4	Parameter-Efficient Fine-Tuning	23
3	Learning and Generalization of Vision Transformers	25
3.1	Introduction	25
3.2	Background and Related Work	26
3.2.1	Problem Formulation and Learning Algorithm	27
3.3	Theoretical Characterization of Transformers	29

3.3.1	Main Theoretical Insights	29
3.3.2	Data Model.....	30
3.3.3	Formal Theoretical Results.....	31
3.4	Performance Evaluation.....	33
4	Formalizing In-Context Learning in Transformers.....	37
4.1	Introduction.....	37
4.2	Background and Related Work	39
4.2.1	Formalizing In-Context Learning with Transformers.....	40
4.3	Theoretical Characterization of In-Context Learning	43
4.3.1	Main Theoretical Insights	43
4.3.2	The Modeling of Training Data and Tasks	43
4.3.3	In-Domain and Out-of-Domain Generalization with Sample Complexity Analysis.....	45
4.3.4	ICL With Magnitude-Based Model Pruning	47
4.3.5	The Mechanism of ICL by the Trained Transformer	48
4.4	Performance Evaluation.....	50
 Part II Advanced Topics in Foundation Models		
5	Automated Visual Prompting	55
5.1	Introduction.....	55
5.2	Background and Related Work	57
5.3	AutoVP Framework	59
5.3.1	Input Scaling	60
5.3.2	Visual Prompt.....	60
5.3.3	Pre-trained Classifier	61
5.3.4	Output Label Mapping	61
5.3.5	End-to-End Hyper-Parameter Tuning	63
5.4	Performance Evaluation.....	64
6	Prompting Large Language Models with Privacy	67
6.1	Introduction.....	67
6.2	Background and Related Work	68
6.2.1	Differential Privacy (DP)	68
6.2.2	Document Processing with Privacy	70
6.3	DP-Prompt.....	71
6.4	Performance Evaluation.....	73
7	Memory-Efficient Fine-Tuning for Foundation Models	77
7.1	Introduction.....	77
7.2	Algorithmic Foundations of ZO Optimization	78
7.3	Applying ZO Optimization for Memory-Efficient Fine-Tuning ...	81
8	Large Language Models Meet Time Series	87
8.1	Introduction.....	87
8.2	Background and Related Work	89

8.3	Time-LLM	90
8.3.1	Model Structure	91
8.4	Performance Evaluation	94
9	Large Language Models Meet Speech Recognition	99
9.1	Introduction	99
9.2	Background and Related Work	101
9.2.1	ASR Rescoring and Error Correction	101
9.2.2	Noise-Robust ASR	102
9.2.3	HyPoradise (HP) Benchmarks	102
9.3	Noise-Aware Generative Error Correction	103
9.3.1	Language-Space Noise Embedding	104
9.3.2	Audio Noise Distillation	105
9.4	Performance Evaluation	106
10	Benchmarking Foundation Models Using Synthetic Datasets	109
10.1	Introduction	109
10.2	Background and Related Work	112
10.3	SynBench	113
10.3.1	Synthetic Data	113
10.3.2	Main Theorem	114
10.3.3	Objective	116
10.3.4	Robustness-Accuracy Quantification	117
10.4	Performance Evaluation	118
11	Machine Unlearning for Foundation Models	123
11.1	Introduction	123
11.2	Research Objective, Formulation, and Related Work	124
11.3	Sparse Optimization for MU: Leveraging Model Sparsity for Efficient and Effective Unlearning	129
11.4	Second-Order Optimization for MU: Iterative Influence-Guided Unlearning	134
11.5	Adversarial Evaluation of MU	137

Part III Trust and Safety in Foundation Models

12	Trustworthiness Evaluation of Large Language Models	149
12.1	Introduction	149
12.2	Background and Related Work	152
12.2.1	Large Language Models (LLMs)	152
12.2.2	Evaluation on LLMs	154
12.2.3	Trustworthiness-Related Benchmarks	155
12.3	Guidelines and Principles for Trustworthiness Assessment of LLMs	155
12.3.1	Truthfulness	156
12.3.2	Safety	157
12.3.3	Fairness	158

12.3.4	Robustness	158
12.3.5	Privacy	159
12.3.6	Machine Ethics	160
12.3.7	Transparency	160
12.3.8	Accountability	161
12.4	Main Insights from TrustLLM Evaluation	162
12.4.1	Overall Observations	162
12.4.2	Novel Insights into Individual Dimensions of Trustworthiness	164
13	Attacks and Defenses on Aligned Large Language Models	167
13.1	Introduction	167
13.2	Background and Related Work	169
13.3	Gradient Cuff	170
13.3.1	Refusal Loss Function and Landscape Exploration	170
13.3.2	Gradient Norm Estimation	172
13.3.3	Gradient Cuff: Two-Step Jailbreak Detection	173
13.3.4	Performance Evaluation	174
13.4	Defensive Prompt Patch	176
13.4.1	Preliminaries	176
13.4.2	Score Evaluation	177
13.4.3	DPP Training Algorithm	178
13.4.4	Performance Evaluation	182
14	Safety Risks in Fine-Tuning Large Language Models	185
14.1	Introduction	185
14.2	Background and Related Work	187
14.3	Performance Evaluation	189
14.3.1	Experiment Setup	189
14.3.2	Numerical Results	190
15	Watermarks for Large Language Models	195
15.1	Introduction	195
15.2	Background and Related Work	196
15.3	Duwak: Dual Watermarking for LLMs	198
15.3.1	Token Probability Watermark	199
15.3.2	Contrastive Search Watermark	199
15.3.3	Watermark Detection in Duwak	200
15.4	Performance Evaluation	202
15.4.1	Experiment Setup	202
15.4.2	Numerical Results	203
16	AI-Generated Text Detection	209
16.1	Introduction	209
16.2	Background and Related Work	210
16.3	RADAR: Robust AI-Text Detection Using Adversarial Learning	212

16.3.1	Training Paraphraser via Clipped PPO with Entropy Penalty	213
16.3.2	Training Detector via Reweighted Logistic Loss	214
16.3.3	RADAR Algorithm	215
16.4	Performance Evaluation	215
16.4.1	Experiment Setup	215
16.4.2	Performance Evaluation and Comparison with Existing Methods	218
16.4.3	AI-Text Detection Transferability of RADAR	219
17	Backdoor Risks in Diffusion Models	221
17.1	Introduction	221
17.2	Background and Related Work	223
17.3	VillanDiffusion: A Unified Backdoor Attack Framework	224
17.3.1	Backdoor Unconditional Diffusion Models as a Distribution Mapping Problem	224
17.3.2	Generalization to Various Schedulers	225
17.3.3	Generalization to ODE and SDE Samplers	227
17.3.4	Unified Loss Function for Unconditional Generation with Image Triggers	228
17.3.5	Generalization to Conditional Generation	229
17.4	Backdoor Detection and Mitigation for Diffusion Models	230
17.5	Performance Evaluation	231
17.5.1	Experiment Setup	231
17.5.2	Caption-Trigger Backdoor Attacks on Text-to-Image DMs	232
17.5.3	Image-Trigger Backdoor Attacks on Unconditional DMs	234
18	Prompt Engineering for Safety Red-Teaming: A Case Study on Text-to-Image Diffusion Models	237
18.1	Introduction	237
18.2	Background and Related Work	238
18.3	Prompting4Debugging (P4D)	241
18.4	Ring-A-Bell	243
18.5	Performance Evaluation	245
18.5.1	P4D Results	246
18.5.2	Ring-A-Bell Results	249
	References	255
	Index	309

Part I

Fundamentals of Foundation Models

Chapter 1

Foundation Models and Generative AI



Abstract This chapter begins by defining the scope of foundation models and generative AI and discussing their differences. We then highlight the unique machine learning paradigms by introducing the neural scaling laws and emerging capabilities of foundation models. We explore the landscape of foundation models through the lens of the AI lifecycle and common practices, and present some canonical examples of foundation models. Finally, we provide our perspective on foundation models.

1.1 What Are Foundation Models and Generative AI?

Foundation model is a technical term coined by Bommasani et al. [67] to highlight a significant paradigm shift in machine learning. Without loss of generality, foundation models are high-capacity neural networks (e.g., neural networks with billions of trainable parameters) trained with large-scale data (e.g., the entire text data scraped from the Internet). Once a foundation model is trained, it can be used to solve various downstream machine learning tasks. Some examples of the data modalities and downstream tasks of foundation models is illustrated in Fig. 1.1. While the training and tuning of foundation models are costly in time and resources, this “one-for-all” methodology deviates from the conventional “one-for-one” principle that trains one specific model for one task. For example, convolutional neural networks (CNNs) are often used in vision tasks such as image recognition or object detection, whereas long short-term memory (LSTM) models are often used in natural language processing tasks such as sentiment classification or summarization. Foundation models change the landscape of machine learning research and technology by sparing the need for training task-specific models, thereby making a unified foundation for different tasks. Specifically, with the advent of transformer architecture [844] featuring attention mechanisms, foundation models can be easily applied to different data modalities (e.g., text, image, and audio, etc) by converting the raw data into some form of discrete “tokens” (mostly a sequence of tokens). A tokenizer could be viewed as a dictionary of tokens (e.g., sub-words in text) of a data modality. Tokenization describes the procedure to




Fig. 1.1 Overview of foundation models: data modalities and downstream tasks

represent a piece of information (e.g., a sentence) as a sequence of tokens. Moreover, transformers also facilitate the integration of multiple modalities due to similar architectural design and tokenization, giving rise to a plethora of multi-modal foundation models. We will highlight some popular single-modal and multi-modal foundation models in Sect. 1.4. The details of the neural networks behind the foundation models are provided in Chap. 2.

On the other hand, *generative AI* (GenAI) refers to machine learning technology and applications featuring generative capabilities, such as large language models (LLMs) that can understand human users' instructions (also known as the "context") and generate the desired content and actions. For instance, one can ask an LLM to read a document (input to the LLM) and make a summary (output from the LLM). However, it is worth noting the discrepancy between foundation models and GenAI. As illustrated in Fig. 1.2, while it is true that many GenAI applications (e.g., LLMs) are empowered by foundation models (the part where they overlap), foundation models and GenAI are not mutually inclusive. Specifically, some foundation models serve as generic data encoders that encode data samples into numerical vectors (or tensors) also known as *embeddings* or *representations*. These data representations provided by foundation models are then further processed by other machine learning models (e.g., a linear classifier or a K-nearest neighbor algorithm) to adapt to a downstream task. Therefore, in this case, foundation models are used for non-generative tasks. On the other hand, many generative AI applications may not necessarily rely on foundation models. For example, in forecasting tasks that predict future events, one can use statistical prediction tools that are not based on deep learning. Furthermore, a chatbot application can be a ruled-based system that generates responses based on predefined states and rules. In this book, we will focus on the scenarios of using foundation models as encoders, as well as generative AI applications based on foundation models.




Fig. 1.2 Comparison of foundation models v.s. generative AI applications using the Venn diagram

We also want to remark that the exact definitions and context of “foundation models” and “GenAI” will likely evolve with technological advances. Looking back, AI’s scope and mainstream technology are dynamically changing and often expansive over time. Similarly, what is considered a foundation model today (e.g., an autoregressive LLM with 7B parameters) might be obsolete if there are more capable and efficient models, i.e., the next-generation foundation models. Consequently, this book takes a broader and more persistent view that focuses on the technical methodologies and fundamental challenges in foundation models and GenAI, instead of taking a narrower view defining what they are and what they can do. Unless otherwise noted, this book positions foundation models as the backbone technology that supports generative and non-generative machine learning tasks.

1.2 Foundations Models at Scale

1.2.1 Neural Scaling Laws

Neural networks are the key ingredients in deep learning techniques [440]. For many deep learning based models, scaling up different factors associated with model training is a practically effective strategy to improve the model’s performance and desirable properties. The study of neural network scalability with respect to a certain

performance metric or property (e.g., test loss of downstream task performance) is known as the neural scaling laws [306, 385]. The scalability analysis often includes

- *Data (tokens)*: How does the model performance scale with the size of training data samples (similarly, the number of tokens)? The study can also include data quality, diversity, and synthetic data generated by a foundation model.
- *Model size (parameters)*: How does the model performance scale with the number of trainable parameters in the neural network model? Specifically, given one type of neural networks (e.g., transformers v.s. CNNs), which type scales better with the model size?
- *Compute*: How does the model performance scale with the level of computing capabilities for training, such as the scaling law with respect to the floating-point operations per second (FLOPS)?

An interesting observation from the neural scaling laws of foundation models is that in most current studies, there is no obvious sign of performance saturation or diminishing returns as training resources are scaled up to the available resources. This means that while foundation models can be quite resource-intensive, based on current scaling laws, one can reasonably expect improved performance at the cost of increased data volume, model size, and computation, provided that these resources have not been exhausted. Nonetheless, new approaches to improve the scaling laws (equivalently, the training efficiency), such as higher-quality data, better model architecture design and training algorithms, as well as more advanced computing hardware, are active and important research directions.

1.2.2 Emerging Abilities

In the study of foundation models, neural scaling laws are often accompanied by the discussion on “emerging abilities” when scaling up the data, model size, or compute power. For example, given the same training data, once the backbone model exceeds a certain model size, the trained model is said to acquire certain abilities that won’t be observed in smaller models [887]. More generally, emergence refers to the capabilities of foundation models that appear suddenly and unpredictably while scaling up. This observation also raises an interesting question of whether one can reliably predict the capabilities of foundation models if a model can acquire some abilities unexpectedly. Sometimes, the acquired abilities can generalize beyond what the backbone model is trained for. For example, an autoregressive LLM is trained to maximize the log-likelihood for the next-token prediction task. However, many studies have found some surprising emerging abilities of such LLMs when scaling up the model size, such as the ability in some level of reasoning [887], where reasoning is a new property to the training objective of next-token prediction. Notably, in [728], the authors argue that the observed “emergence” (mathematically, an abrupt change in a metric when scaling up) can be an artifact of a chosen metric

for performance measurement. The effect of emergence might be less obvious if one chooses a linear or continuous metric. Nonetheless, the scaling properties of a foundation model remain fundamentally important, regardless of the existence of emergent abilities.

Overall, scalability plays an important role in foundation models. It also shifts the research focus from “can a model solve such problems?” to “what problems can be properly addressed by scaling up?” versus “what problems are persistent after scaling up?” Furthermore, what are the new challenges that could arise after scaling up? Many ongoing discussions are made around whether scalability and emerging abilities are the paths toward building machine learning systems surpassing human-level intelligence (which is one of the many ways to define artificial general intelligence, AGI), as well as causing catastrophic risks threatening our safety, the society, and the environment.

1.3 Lifecycle of Foundation Models

Following the notion of AI lifecycle in [121], here we define the lifecycle specific to foundation models. In general, an AI lifecycle can be divided into two phases: *development* and *deployment*. The development phase contains two stages: *data preparation* and *model training*. Figure 1.3 depicts the 3-stage lifecycle for foundation models, including (i) data preparation, (ii) model training, and (iii) model deployment. In each stage, we also highlight some key components and common practices. Note that the lifecycle is recurrent. Initially, one starts with preparing data and model training for the first deployment. After deployment, the model can return




Fig. 1.3 Lifecycle of foundation models

to the data preparation and model training stages if further updates and refinements are required. In what follows, we provide a high-level overview of each stage and the associated components in Fig. 1.3. We defer their details to later sections.

1.3.1 Data Preparation

- *Data Collection* is the process of gathering training and evaluation data for foundation models. When pre-training a foundation model to learn generic data representations (e.g., self-supervised training through next-token prediction), a common practice is to collect large-volume but possibly noisy data, such as all textual data scraped from the Internet. Data collection also involves gathering high-quality and expert-annotated data samples for efficient fine-tuning.
- *Data Pre-processing* describes the procedure to convert the collected data into a unified format (e.g., tokenization) and the data cleansing steps to improve the quality of the collected data. To improve the trustworthiness of foundation models, data pre-processing often involves extra steps to detect and remove undesirable content, such as bias, toxicity, private data, copy-righted data, etc.
- *Data Augmentation* refers to methods to expand the data collection. For example, augmenting noisy or semantic-preserving image transformations is an effective strategy to improve image foundation models. Collecting synthetic data generated by other (often large) foundation models, as a form of knowledge distillation, is also a popular approach to train a smaller but capable language model. Moreover, after the model is deployed, any failure modes and mistakes made by the deployed model can also be collected for future training.

1.3.2 Model Training

- *Pre-training* often features self-supervised learning over a dataset. Self-supervised learning means the model parameters are trained using some self-defined objectives on the data samples and without the supervision of any additional annotations (e.g., labels made by humans) [253], which makes pre-training scalable to large datasets. Supervised pre-training is plausible if some supervised information is available. For image data, self-supervised learning can be realized by learning to align the representations of the original and corrupted (e.g., random masking) image pairs, such as contrastive learning [129] and masked autoencoders [289]. For sequential data such as text and temporal signals, self-supervised learning can be realized by masking some tokens in a sequence and training the model to predict the masked tokens. For example, a base autoregressive LLM is usually pre-trained with a self-supervised auto-completion task to predict the next token.

- *Instruction-tuning* finetunes a pretrained model to follow instructions. A typical example is fine-tuning a base LLM pre-trained for autocompletion on a labeled dataset of instructional prompts and corresponding outputs to teach an LLM to communicate like a human, such as a conversational agent or a chat model. Instruction-tuning is often realized by supervised fine-tuning on labeled input-output pairs to accommodate the response of foundations to follow desirable forms, such as answering a user's question with a list of points or steps.
- *Alignment* is usually jointly considered with instruction-tuning to align the output of the foundation models to human values. Specifically, safety alignment refers to the effort to reduce the risks (or to enhance the safety guardrails) by fine-tuning a foundation model to prevent harm and follow compliance and regulations. More broadly, alignment also includes adapting foundation models to cultural norms, addressing privacy and fairness concerns, mitigating security vulnerabilities, improving interpretability and transparency, and red-teaming foundation models. Common instruction-tuning and alignment techniques include supervised fine-tuning, reinforcement learning with human feedback (RLHF) [625], and direct preference optimization (DPO) [682]. An instruction-tuned model with alignment embedded is called an aligned model.

1.3.3 Model Deployment

Model deployment of foundation models can be further divided into two categories, depending on whether the application scenarios involve updating the model parameters (training-based) or not (training-free).

The training-free model deployment refers to prompt (input) engineering techniques to maximize the utility of a foundation model without updating its trained parameters, which include

- *Zero-shot Prompting* refers to using a trained foundation model by directly providing the user query. For example, asking an LLM to recommend a book on foundation models, or showing an image to a vision-language foundation model to ask questions related to the image content. It is called zero-shot because no examples are given to a foundation model.
- *In-context Learning (ICL)* means providing a set of demonstrations to a foundation model to better understand the context before asking the model to solve a task. For example, one can provide some question-answer pairs as demonstrations, with the questions in one language and the answers in another. Then, after providing the demonstrations, ask the real question so that a foundation model can understand the expected answer. Few-shot ICL means a small number of examples are provided, along with the user query, to prompt a foundation model.
- *Chain-of-Thought (CoT)* refers to effective methods to prompt a foundation model to further reason about the user query and provide more constructive and accurate responses [888]. One typical example is to simply append "Let's think

step-by-step” to the user query, which is empirically shown to be effective in improving the performance of downstream tasks by helping foundation models break a complex problem into a series of sub-problems or steps that are easier to tackle. CoT can also be used together with ICL, such as few-shot ICL with CoT (few-shot CoT).

The training-based model deployment refers to efficient fine-tuning approaches to adapt a trained foundation model, which include

- *Parameter-efficient fine-tuning (PEFT)* is a cost-efficient approach to adjust a trained foundation model by inserting a relatively small number of trainable parameters into a pre-trained model. Popular examples include pre-fix tuning [483], prompt tuning (by adding some additional trainable tokens, also known as soft prompts) [452], adapters, and low-rank adaptation (LoRA) [322].
- *Full Fine-tuning* refers to the fine-tuning process that involves updating all (or most) parameters of trained foundation models. While full fine-tuning can be effective in improving foundation models, the associated training resources, such as memory usage and computing hardware, can also be quite costly.
- *Model Merging* explores the potential of merging foundation models of complementary capabilities, such as merging a chat model with a specialized math foundation model to become a chat model with improved reasoning skills. It also includes the studies of model ensembling [905] and task arithmetic [347] to edit the functionalities of foundation models.

1.4 Canonical Examples of Foundation Models

As the fields of foundation models and generative AI technology are rapidly growing, we introduce some canonical examples of foundation models that have made far-reaching impacts. More foundation models will be introduced in subsequent chapters.

- **Generative Pre-trained Transformer (GPT)** is a transformed-based neural network component that empowers state-of-the-art foundation models, especially for large language models [81]. It is worth noting that GPT stands for the last three letters of ChatGPT, a signature generative AI service via chat interfaces by OpenAI. Transformers feature self-attention mechanisms for learning context from data. GPT models are shown to have good scalability with increased pre-training data volume and model size [385]. The generality of GPT also facilitates the training of foundation models with different modalities, through proper tokenization over multi-modal data.
- **Contrastive Language-Image Pre-training (CLIP)** [674] is a multi-modal encoder model that uses contrastive learning with image-text pairs to train a text encoder and an image encoder through a joint embedding space. CLIP features a revolutionary approach to encoder different data samples of different modalities into a common latent space and facilitates single-modality and multi-modality

downstream tasks. It also inspires the design of other multi-modal foundation models beyond languages and images.

- **Diffusion model** is a state-of-the-art generative model [298, 768], especially in continuous data domains such as image, audio, and video. In principle, diffusion models follow a mathematical diffusion process to gradually encode data samples into random Gaussian noises, and train a neural network to learn to decode (reconstruct) the data samples based on multiple denoising steps. Many diffusion models are further empowered with text-based instructions to generate high-quality realistic data samples adhering to user queries. For example, many text-to-image generation applications use a CLIP model to obtain the context embedding of a user query and pass it to a diffusion model for image generation.

1.5 Our Perspective on Foundation Models

Based on the advances and trends observed in foundation models and GenAI technology, we provide our perspective on their impacts and implications as follows.

- **Foundation models are the new essentials:** Capable foundation models encode raw data into informative representations and facilitate the downstream machine learning tasks. Future research and technology should make use of the data representations as the new foundation. Furthermore, when raw data of different kinds become some form of unified embeddings, efficient techniques for prompts, instructions, and demonstrations to improve the associated task performance are the new data.
- **Governance (risk management) is the new race:** We predict that near-term foundation models (including multi-modal variants) will all become similar in capability as available data resources are exhausted. Instead, a key differentiator is which foundation model has safer and more trustworthy features. More emphasis will be put on the governance of risk management of foundation models.
- **AI research is becoming empirical science:** In the line of foundation models, developers embrace the boldness with an aim to build complex frontier AI systems fast and understand how to safely and responsibly use them later. Also, the community acknowledges imperfection—fast-paced development and deployment are accompanied by notable sociotechnological challenges and risks, with the hope that these issues can be addressed and fixed on-the-fly. Finally, the rigor of adopting scientific methods are expected to ensure safe and sustainable use of foundation models, especially in high-stakes application scenarios.

The goal of this book is to provide a comprehensive overview and technical deep dives of foundation models, so that the readers can gain a fundamental understanding of this rapidly growing research field and be motivated to develop a better, safer, and more trustworthy AI technology.

Chapter 2

Neural Networks



Abstract The chapter begins with the introduction of basic modules in modern neural networks. Then, we provide details about transformers, which are state-of-the-art neural network architectures and popular backbones for foundation models. Finally, we summarize major components in large language models, including next-token prediction, decoding, alignment, and parameter-efficient finetuning.

2.1 Introduction

Without loss of generality, a neural network is a parametrized composition function consisting of a set of layers (or blocks of layers). These layers or blocks are basic modules that can be configured to make a neural network model. The parameters associated with these modules are often updated by some gradient-based optimization algorithms, such as the Adam optimizer [404], where the training objective function is evaluated on a batch of data samples. Features or embeddings are often used to describe the output of a particular layer, as a form of the latent/internal representations of data samples. While there is no precise definition of what makes a “deep” or “large” neural network, the training cost of a neural network, including memory usage and compute time, is associated with the number of parameters and the induced features. In addition, the gradient computation relies on the backpropagation and automatic differentiation capabilities enabled by deep learning packages, which require hardware accelerations such as graphics processing units (GPUs).

Let x denote a data input and let f_θ denote an L -layer neural network model defined as

$$f_\theta(x) = f_{\theta_L}^{(L)} \circ f_{\theta_{L-1}}^{(L-1)} \circ \dots \circ f_{\theta_1}^{(1)}(x), \quad (2.1)$$

where $f_{\theta_\ell}^{(\ell)}$ denotes the ℓ -th layer ($\ell \in \{1, 2, \dots, L\}$), θ_ℓ denotes the parameters of the ℓ -th layer (θ_ℓ can be an empty set, $\theta_\ell = \emptyset$, if the layer does not have any trainable parameters), $\circ$ denotes the mathematical operation of function composition, and $\theta = \{\theta_1, \theta_2, \dots, \theta_L\}$ denotes the collection of all parameters.

Section 2.2 introduces some basic and common neural network modules. Section 2.3 focuses on transformers, as they are popular backbones for many state-of-the-art neural network architectures. Finally, Sect. 2.4 summarizes key components of large language models.

2.2 Basic Neural Network Modules

In this section, we define and summarize some basic neural network modules. There are three popular types of neural network layers: *activation*, *parametrized*, and *pooling*. Activation layers refer to non-parametric and often nonlinear operations on features to improve the expressiveness of data representations. Parametrized layers refer to operations with trainable parameters. Pooling layers refer to operations aiming to reduce the feature size.

For ease of illustration, we will also use x to denote a layer's input, use y to denote a layer's output, and W to denote the trainable parameters of a layer. We also assume $x \in \mathcal{R}^{d_{in}}$ and $y \in \mathcal{R}^{d_{out}}$ are in real numeric vector forms. Extensions to higher-order cases are possible but may not be uniquely defined. For instance, if $x \in \mathcal{R}^{d_{in} \times h}$ and $y \in \mathcal{R}^{d_{out} \times h}$, then W can be designed as a weight-shared operation along each column dimension, or a matrix incorporating $d_{in} \times h \times d_{out} \times h$ parameters, among other possible designs.

- **Linear layer:** Linear layer is also called a multi-layer perception (MLP) or a fully-connected (FC) layer. It is defined as

$$y = Wx + b, \quad (2.2)$$

where b is known as the bias term which is trainable. Nonetheless, one can rewrite the linear layer with a bias term as another equivalent linear transformation $y = W'x'$, where $W' = [W \ b]$ and $x' = \begin{bmatrix} x \\ 1 \end{bmatrix}$.

- **Convolution layer:** Convolution layer is a common module in deep learning models for computer vision, especially for two-dimensional (2D) convolution. One standard form of 2D convolution layer is a $d \times d$ array of shared weights (also known as a kernel or a filter) that moves across an image to extract features by applying the weights to the corresponding local pixel values. The number of filters affects the depth of the output. There are additional hyperparameters, such as stride and padding, that control how the kernel moves and ultimately affect the size of output features. One often uses $y = W * x$ to denote the features extracted from a convolution layer with a kernel W .
- **ReLU layer:** The rectified linear unit (ReLU) is a popular activation function to introduce nonlinearity into features. It is an element-wise operation defined as

$$y = \text{ReLU}(x) = \max\{0, x\}, \quad (2.3)$$

which essentially filters out the negative elements in x and replaces them with a value of 0, while retaining other elements of x .

- **GeLU layer:** Extending from ReLU, the Gaussian Error Linear Unit (GeLU) [294] is widely used as an activation function in transformer-based architectures. It is defined as

$$y = \text{GeLU}(x) = x \cdot \Phi(x), \quad (2.4)$$

where $\Phi(x)$ is the cumulative distribution function of a standard Gaussian random variable. In comparison, GeLU weights the input by its percentile, whereas ReLU truncates the input by its sign.

- **Pooling layer:** Pooling is an effective operation to reduce the dimension of features. It is often a non-parametric operation that does not involve any trainable parameters. Some popular choices are average-pooling, which takes the averaged input values as the output, and max-pooling, which takes the maximum input value as the output.
- **Softmax layer:** softmax is a nonlinear operation that maps features to non-negative values summing to 1. In neural network classification models, it has been the default choice to generate confidence scores (prediction probabilities) for each class. In addition, it is also used in the attention layer of transformers. The i -th element of the output y in the softmax layer is defined as

$$y_i = \text{softmax}(x) = \frac{\exp(x_i/T)}{\sum_{k=1}^{d_{in}} \exp(x_k/T)}, \quad (2.5)$$

where T is called the *temperature*, a hyperparameter affecting the nonlinearity of the output. If $T \rightarrow \infty$, the output y will have a uniform distribution across each dimension. On the other hand, if $T \rightarrow 0$, y will have a skewed and concentrated distribution on the most probable class. The input (x) to the softmax function is also called the *logits*.

- **Residual block:** Residual block contains a set of layers whose output is a linear combination of the last layer and the input of some preceding layer [290]. Such a combination is enabled by *skip connection*. Let x denote the block input and let $\{W^{(k)}\}$ denote the parameters associated with the layers involved in the block. Then, let $F(x, \{W^{(k)}\})$ denote the output of the block without skip connection. The final output of the residual block can be defined as

$$y = F(x, \{W^{(k)}\}) + W_s x, \quad (2.6)$$

where W_s is a linear projection that matches the input dimension to the output. If y and x have the same dimension, then W_s can be omitted.

- **Batch normalization layer:** Batch normalization (BN) [351] is a practical operation to normalize the layer input and accelerate neural network training. When training neural networks with a mini-batch, batch normalization computes

the mean and variance of the i -th feature based on the mini-batch, denoted as m and σ^2 , respectively, and obtains a normalized feature (of the i -th element) as

$$\hat{x}_i = \text{BN}(x_i) = \frac{x_i - m}{\sqrt{\sigma^2 + \epsilon}}, \quad (2.7)$$

where $\epsilon > 0$ is some small positive constant to avoid numerical instability. Finally, to restore the representation power of the neural network, batch normalization adds a trainable linear transformation on the normalized feature and expresses the final output as

$$y_i = \gamma \cdot \hat{x}_i + \beta, \quad (2.8)$$

where γ and β are trainable parameters. Note that after training, the model uses the mean and variance of the entire training data batch in the batch normalization layer for inference.

2.3 Transformers

Transformers are state-of-the-art neural network architectures for the representation learning of sequence data. We follow the mathematical framework in [653] to explain transformers. We assume the maximum input length (number of tokens) to a transformer is $\ell_{\max}$. For ease of presentation, we may abuse the notation of row vectors and column vectors to avoid using transpose $\cdot^T$ to overcomplicate the formulation. In this part, we use bold mathematical notation to denote vectors or matrices, unbold mathematical notation for scalars, and calligraphic mathematical notation for sets.

2.3.1 Token Embedding and Position Encoding

Token Embedding Let $\mathcal{V}$ denote the vocabulary of tokens and let $[N_V] := \{1, 2, \dots, N_V\}$ denote the index set of tokens, where $|\mathcal{V}| = N_V$. Let $\mathbf{x}$ denote a sequence of tokens, where $\mathbf{x}[i] \in [N_V]$ denote the index of the i -th token (i.e., the token ID) in the vocabulary $\mathcal{V}$. Each token in $\mathcal{V}$ is represented by a numerical vector $\mathbf{e}$ of dimension d_e (d_e is also called the embedding size). An embedding matrix $\mathbf{H} \in \mathcal{R}^{N_V \times d_e}$ is constructed by stacking all token embeddings row-wise, where $\mathbf{h}_i$, the i -th row of $\mathbf{H}$, is the token embedding of the i -th token in $\mathcal{V}$. A trainable token embedding can be parametrized by a linear weight matrix $\mathbf{W} \in \mathcal{R}^{d_e \times N_V}$.

Position Encoding Position encoding aims to represent the order of tokens in a sequence as an embedding. It can be either learned or hard-coded. The original

transformer paper [844] uses a sinusoidal function for position embedding of text. A trainable position encoding can be parametrized by a linear weight matrix $\mathbf{W} \in \mathcal{R}^{d_e \times \ell_{\max}}$. The form of position embedding can vary, depending on the data characteristics. For example, vision transformers can use spatial (two-dimensional) position encoding when treating image patches as tokens [193]. Let $\mathbf{p}_j \in \mathcal{R}^{d_e}$ denote the position embedding vector of the j -th token in $\mathbf{x}$ ($1 \leq j \leq \ell_{\max}$). Recall $\mathbf{h}_{\mathbf{x}[j]}$ is the token embedding vector of the token ID $\mathbf{x}[j]$. Then, the final embedding of such token is $\mathbf{e} = \mathbf{h}_{\mathbf{x}[j]} + \mathbf{p}_j$.

2.3.2 Attention

Attention Attention is a critical design in transformers. Let $\mathbf{x} = \{\mathbf{x}_j\}$ denote input token embeddings and let $\mathbf{y} = \{\mathbf{y}_i\}$ denote the output token embeddings after attention. The embedding of the i -th token in $\mathbf{y}$ is represented as $\mathbf{y}_i = \sum_{j \in \mathcal{J}} a_{i,j} \mathbf{x}_j$, where $a_{i,j}$ denotes the attention weight of $\mathbf{y}_i$ on the j -th token in $\mathbf{x}$ (i.e., $\mathbf{x}_j$), and $\mathcal{J}$ denotes the set of tokens in $\mathbf{x}$ contributing to the token embedding of $\mathbf{y}_i$ through attention weights. In practice, the attention weights $\{a_{i,j}\}$ are realized through a softmax function, which will be detailed later in this section. Self-attention means the attention operation is implemented only on the same input sequence. In other words, the output sequence is a weighted representation (via nonlinear aggregation) of the input sequence through self-attention. Cross-attention means the attention operation is implemented by combining or mixing two or more different input sequences (e.g., fusing token embeddings of multiple modalities). If the attention is bidirectional (or fully visible), then $\mathcal{J}$ contains all tokens in $\mathbf{x}$. If the attention is unidirectional (or causal), then $\mathcal{J}$ contains all tokens preceding the token j and includes j itself. That is, $\mathcal{J} = \{j : j \leq i\}$. Unidirectional attention masks out future tokens and enables autoregressive prediction for the next tokens given the current context (i.e., observable tokens and their embeddings).

Query, Key, and Value In addition to attention weights, the attention mechanism in transformers also consists of three main components: *query*, *key*, and *value*, which are trainable linear transformations of input token embeddings. These linear transformations follow the form of a linear projection matrix ($\mathbf{W}$) on the token embeddings plus a bias term ($\mathbf{b}$). Let $\mathbf{1}_k$ denote a column vector of ones with size k . Consider two transposed token embedding matrices $\mathbf{X} \in \mathcal{R}^{d_x \times \ell_x}$ (the primary sequence) and $\mathbf{Z} \in \mathcal{R}^{d_z \times \ell_z}$ (the context sequence), where d_x/d_z are their token embedding sizes and ℓ_x/ℓ_z are their token lengths. The attention mechanism on $\mathbf{X}$ and $\mathbf{Z}$ will produce an output token embedding matrix $\mathbf{Y} \in \mathcal{R}^{d_y \times \ell_y}$ and $\ell_y = \ell_x$ (the output token length is the same as the primary sequence). The Query ($\mathbf{Q}$), Key ($\mathbf{K}$), Value ($\mathbf{V}$) matrices are defined as

$$\mathbf{Q} = \mathbf{W}_q \mathbf{X} + \mathbf{b}_q \mathbf{1}_{\ell_x}^T; \quad \mathbf{K} = \mathbf{W}_k \mathbf{Z} + \mathbf{b}_k \mathbf{1}_{\ell_z}^T; \quad \mathbf{V} = \mathbf{W}_v \mathbf{Z} + \mathbf{b}_v \mathbf{1}_{\ell_z}^T. \quad (2.9)$$

where $\mathbf{W}_q \in \mathcal{R}^{d_a \times d_x}$, $\mathbf{b}_q \in \mathcal{R}^{d_a}$, $\mathbf{W}_k \in \mathcal{R}^{d_a \times d_z}$, $\mathbf{b}_k \in \mathcal{R}^{d_a}$, $\mathbf{W}_v \in \mathcal{R}^{d_y \times d_z}$, $\mathbf{b}_v \in \mathcal{R}^{d_y}$, d_a is the latent dimension of the linear projection in attention, and d_y is the token embedding size after attention. The dimensions of $\mathbf{Q}$, $\mathbf{K}$, $\mathbf{V}$ are $\mathbf{Q} \in \mathcal{R}^{d_a \times \ell_x}$, $\mathbf{K} \in \mathcal{R}^{d_a \times \ell_z}$, and $\mathbf{V} \in \mathcal{R}^{d_y \times \ell_z}$, respectively. Finally, the output embedding matrix is defined as

$$\mathbf{Y} = \mathbf{V} \cdot \text{softmax} \left(\frac{\mathbf{K}^T \mathbf{Q}}{\sqrt{d_a}} \right). \quad (2.10)$$

Here, $\mathbf{V}$ and $\text{softmax} \left(\frac{\mathbf{K}^T \mathbf{Q}}{\sqrt{d_a}} \right)$ are both matrices, and $\cdot$ means matrix-matrix product. Equation (2.10) can be viewed as a matrix representation of the attention weights operated on each column (token) of the columns in $\mathbf{V}$, where the attention weights for each column are realized through the softmax function, $\text{softmax}(\cdot)$, which is defined element-wise as $\text{softmax}(\mathbf{A})[t_z, t_x] := \frac{\exp \mathbf{A}[t_z, t_x]}{\sum_{t \in \mathcal{J}_t} \exp \mathbf{A}[t, t_x]}$, where t_z/t_x are the row/column indices of the matrix $\mathbf{A}$, and $\mathcal{J}_t$ describes the set of tokens tacking effect on the attention weight of the token index t . For self-attention, the attention reduces to $\mathbf{X} = \mathbf{Z}$.

We provide more interpretations of the attention function in (2.10) as follows.

- The matrix-matrix product $\mathbf{K}^T \mathbf{Q}$ inside the softmax function is called the attention score matrix, which is an $\ell_z \times \ell_x$ matrix describing the attention weights between the tokens of the two sequences $\mathbf{X}$ and $\mathbf{Z}$. The scaling term $\sqrt{d_a}$ in the denominator is to normalize the matrix-matrix product to prevent the softmax operation from over-saturation.
- Each column of $\mathbf{Y}$ is the token embedding obtained by employing the defined non-linear softmax attention on the columns of the value matrix $\mathbf{V}$.
- By construction, computing the attention function in (2.10) has quadratic computation complexity in the order of $O(\ell_z \cdot \ell_x \cdot d_a)$, meaning that the cost of computing the attention grows quadratically with the length of the input tokens, which is known to be a computation bottleneck of transformers.

Multi-Head Attention Equation (2.10) describes the attention mechanism with a single *attention head* (i.e., only with one set of QKV components). In practice, *multi-head attention* is used in transformers, which runs multiple heads (separable sets of QKV components) on the same input in parallel and combine their outputs through concatenation. Assume there are H heads, each resulting in an output embedding matrix $\mathbf{Y}^{(h)} \in \mathcal{R}^{d_y \times \ell_y}$, $h \in \{1, 2, \dots, H\}$. The concatenated embedding

matrix from multi-head attention is represented as $\mathbf{Y} = \begin{bmatrix} \mathbf{Y}^{(1)} \\ \mathbf{Y}^{(2)} \\ \vdots \\ \mathbf{Y}^{(H)} \end{bmatrix}$, where $\mathbf{Y} \in$

$\mathcal{R}^{H \cdot d_y \times \ell_y}$. To control the output embedding dimension to be d_o , an additional linear projection with linear weight matrix $\mathbf{W}_o \in \mathcal{R}^{d_o \times H \cdot d_y}$ and bias term $\mathbf{b}_o \in \mathcal{R}^{d_o}$ is

applied to $\mathbf{Y}$, such that the final token embedding matrix becomes $\mathbf{Y}_o = \mathbf{W}_o \mathbf{Y} + \mathbf{b}_o \mathbf{1}_{\ell_y}^T$, where $\mathbf{Y}_o \in \mathcal{R}^{d_o \times \ell_y}$.

Layer Normalization Transformers often include a layer normalization function, similar to batch normalization in (2.7), on the activations.

2.3.3 Transformer Types

With the basic neural network components in Sect. 2.2 and the attention components in Sect. 2.3.2, we summarize the major types of transformers.

- *Encoder-decoder* transformer is a sequence-to-sequence attention-based model [844]. The model encodes an input sequence as embeddings through tokenization and attention, and then decodes from the same embeddings to produce the output sequence. The training involves pairs of input-output sequences for learning sequence-to-sequence translation tasks.
- *Encoder-only* transformer means an encoder-decoder model without the decoder part (hence the named “encoder-only”). The training only involves input sequences. BERT (Bidirectional Encoder Representations from Transformers) [182] is a typical example of an encoder-only model, which is often used for language understanding tasks. BERT uses the GeLU activation function in (2.4).
- *Decoder-only* transformer, such as the family of generative pre-trained transformer (GPT), including GPT-2 [680], GPT-3 [81], etc, is the backbone of autoregressive language models. GPT is pre-trained with the next-token prediction task (see Sect. 2.4 for details). Its training involves input sequences with self-attentions and unidirectional (causal) masks.

Transformer Block A decoder-only transformer model essentially contains a series of transformer blocks, followed by a softmax function with a trainable linear transformation matrix $\mathbf{W}_s \in \mathcal{R}^{N_V \times d}$ to map the token embeddings of the final transformer block (with embedding size d) to the probability distribution over all possible N_V tokens in the vocabulary list for decoding. In its simplest form, a transformer block contains the following components sequentially. We use $\mathbf{X} \in \mathcal{R}^{d_x \times \ell_x}$ to denote the input (transposed) embedding matrix to a transformer block. The first block’s input is obtained through token embedding and position encoding described in Sect. 2.3.1.

1. Layer normalization (layer-norm) following (2.7) (operated on each token): $\mathbf{X} \leftarrow \text{layer-norm}(\mathbf{X})$.
2. Multi-head attention (MA-Attention, including query, key, value, and output projection) plus residual connection: $\mathbf{X} \leftarrow \mathbf{X} + \text{MA-Attention}(\mathbf{X})$.
3. Another layer normalization following (2.7) (operated on each token): $\mathbf{X} \leftarrow \text{layer-norm}(\mathbf{X})$.

4. A two-layer fully-connected network (parametrized by the linear weights $\{\mathbf{W}_1, \mathbf{W}_2\}$ and bias terms $\{\mathbf{b}_1, \mathbf{b}_2\}$) with GeLU activation (see (2.4)) and residual connection: $\mathbf{X} \leftarrow \mathbf{X} + \mathbf{W}_2 \text{GeLU}(\mathbf{W}_1 \mathbf{X} + \mathbf{b}_1 \mathbf{1}^T) + \mathbf{b}_2 \mathbf{1}^T$. The dimension of the vector of ones, $\mathbf{1}$, is determined by the column dimension of the corresponding matrix for addition.

Finally, a layer normalization is applied on the output of the final transformer block, $\mathbf{X} \leftarrow \text{layer-norm}(\mathbf{X})$. The next-token prediction probability matrix $\mathbf{P} \in [0, 1]^{N_V \times \ell_x}$, with the t -th column of $\mathbf{P}$ representing the probability of the next token x_{t+1} given the current context (all preceding tokens, including the t -th one), is represented as $\mathbf{P} = \text{softmax}(\mathbf{W}_s \mathbf{X})$, where the softmax operation for a matrix is defined in Sect. 2.3.2.

2.4 Large Language Models

In this section, we introduce key components in autoregressive large language models (LLMs), including next-token prediction, decoding, alignment, and parameter-efficient fine-tuning.

2.4.1 Next-Token Prediction

Given an input context $\mathbf{x}$ (e.g., a user query), an autoregressive language model generates an output sequence $\mathbf{y} = \{y_1, y_2, \dots, y_T\}$ by predicting one token at a time, based on previously generated tokens. Let θ denote the parameters of the LM. The conditional probability of generating $\mathbf{y}$ given $\mathbf{x}$ is modeled as

$$p_\theta(\mathbf{y}|\mathbf{x}) = \prod_{t=1}^T p_\theta(y_t|\mathbf{x}, \mathbf{y}_{<t}), \quad (2.11)$$

where $\mathbf{y}_{<t}$ denotes the generated sequence, $\mathbf{y}_{<t} = \{y_1, y_2, \dots, y_{t-1}\}$, preceding the t -th token, and $\mathbf{y}_{<1} = \emptyset$ (an empty set). $p_\theta(y_t|\mathbf{x}, \mathbf{y}_{<t})$ is parametrized by a softmax function with temperature T , determining the probability of sampling a token y_t in the vocabulary set $\mathcal{V}$. If the temperature T is set to 0, then the output will be deterministic, known as greedy decoding. Instead of using conditional probability, one often uses the following log-likelihood in a training objective function:

$$\log p_\theta(\mathbf{y}|\mathbf{x}) = \sum_{t=1}^T \log p_\theta(y_t|\mathbf{x}, \mathbf{y}_{<t}). \quad (2.12)$$

Next-token prediction is commonly used for pre-training and fine-tuning LLMs. For LLM pre-training on a dataset $\mathcal{D}$, let $\mathbf{y} \sim \mathcal{D}$ denote a sentence of token length $|\mathbf{y}|$ randomly drawn from $\mathcal{D}$. The next-token prediction loss is defined as the expected

cross entropy loss between the actual and predicted tokens:

$$\text{Loss} = -\mathbb{E}_{\mathbf{y} \sim \mathcal{D}} \frac{1}{|\mathbf{y}|} \sum_{t=1}^{|\mathbf{y}|} \log p_{\theta}(y_t | \mathbf{y}_{<t}), \quad (2.13)$$

where $\mathbb{E}_{\mathbf{y} \sim \mathcal{D}}$ denotes expectation over all random drawn samples from $\mathcal{D}$. This loss is also known as *causal language modeling*, by using previous tokens to predict the next token. Moreover, dropping the negative sign, the training objective is equivalent to maximizing the log-likelihood of next-token prediction. Similarly, for supervised fine-tuning (or instruction-tuning) on a dataset $D = \{(\mathbf{x}, \mathbf{y})\}$, where $\mathbf{x}$ is the input and $\mathbf{y}$ is the desired output, the next token prediction loss becomes

$$\text{Loss} = -\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \frac{1}{|\mathbf{y}|} \sum_{t=1}^{|\mathbf{y}|} \log p_{\theta}(y_t | \mathbf{x}, \mathbf{y}_{<t}), \quad (2.14)$$

Beyond causal language modeling, one can introduce a mask and train a foundation model to predict the masked tokens.

2.4.2 Decoding Strategies

Decoding strategies refer to sampling methods for generating the next token. They vary by the level of randomness, diversity, and fluency in the generated output. We summarize the mainstream approaches as follows.

- *Greedy decoding* refers to sampling the next token with the highest probability. It can be realized by setting the temperature parameter in softmax sampling to be 0, resulting in a deterministic output. Its myopic sampling strategy may lead to suboptimal sentences.
- *Beam search* allows for better diversity and fluency than greedy decoding by retaining some memory buffers (or hypotheses) to track high-probability outputs in previous generations and use these beams to select the next token.
- *Top-K sampling* refers to sampling the next token based on the top-K most probable tokens, with some reweighted probability distribution. One drawback of Top-K sampling is that the number of candidate tokens are fixed during generation.
- *Top-p sampling*, also known as *nucleus sampling*, finds the smallest possible set of tokens whose cumulative probability exceeds the probability p , and then samples the next token from this set with some reweighted probability distribution. Through the probability thresholding mechanism, Top-p sampling allows for dynamic and automated adjustment of the size of the candidate tokens. Current LLMs often use top-p sampling (with p ranging from 0.6 to 0.9) as the default choice.

2.4.3 Alignment Strategies

Alignment refers to the procedures to align the response of LLMs to the desired output. Here we summarize the major techniques in LLM alignment. *Supervised fine-tuning (SFT)* on a pre-trained LLM with a prompt-response dataset $D = \{(\mathbf{x}, \mathbf{y})\}$ is the most straightforward approach for alignment, such as using the loss function defined in (2.13)

Reinforcement learning with human feedback (RLHF) [625] is a popular alignment method, which consists of three major steps:

1. Given the context (e.g., a set of prompts), ask human users to write high-quality responses to create a dataset $\mathcal{D}$. Then, use $\mathcal{D}$ to perform supervised fine-tuning on a pre-trained LLM.
2. Use this fine-tuned model to create a reward model, by asking human users about the generated responses' preferences. The reward model is often obtained by replacing the output layer (the next-token classification layer) with a regression layer and training the model using the annotated preferences.
3. Use the reward model to finetune the supervised fine-tuned model in Step 1 based on proximal policy optimization [737].

For a given context $\mathbf{x}$, the Bradley-Terry model is a logistic function measuring the preference between two responses $\mathbf{y}_1$ versus $\mathbf{y}_2$. Let $r(\mathbf{x}, \mathbf{y})$ denote the reward function of the response $\mathbf{y}$ given the context $\mathbf{x}$. The preference of $\mathbf{y}_1$ over $\mathbf{y}_2$, in terms of the probability of choosing $\mathbf{y}_1$ over $\mathbf{y}_2$, is modelled as:

$$p(\mathbf{y}_1 > \mathbf{y}_2 | \mathbf{x}) = \frac{\exp(r(\mathbf{x}, \mathbf{y}_1))}{\exp(r(\mathbf{x}, \mathbf{y}_1)) + \exp(r(\mathbf{x}, \mathbf{y}_2))} = \frac{1}{1 + \exp(r(\mathbf{x}, \mathbf{y}_2) - r(\mathbf{x}, \mathbf{y}_1))}. \quad (2.15)$$

Let $\mathbb{I}\{\mathbf{y}_1 > \mathbf{y}_2\}$ denote the indicator function such that it is 1 when the event $\mathbf{y}_1 > \mathbf{y}_2$ is true and is 0 otherwise. The corresponding loss function to optimize the reward model to predict the correct preference is

$$\text{Loss}(\mathbf{y}_1, \mathbf{y}_2, \mathbf{x}) = -\mathbb{I}\{\mathbf{y}_1 > \mathbf{y}_2\} \log p(\mathbf{y}_1 > \mathbf{y}_2 | \mathbf{x}) - \mathbb{I}\{\mathbf{y}_2 > \mathbf{y}_1\} \log p(\mathbf{y}_2 > \mathbf{y}_1 | \mathbf{x}). \quad (2.16)$$

Instead of training a reward model for alignment, *preference optimization* [682] proposed to directly use the reward as the loss function to fine-tune the model. Let $\pi(\mathbf{y} | \mathbf{x})$ denote the probability of the fine-tuned model selecting $\mathbf{y}$ as the response given $\mathbf{x}$, and let π_{ref} denote that of the original model (before fine-tuning). In DPO, the preference function is defined as

$$p(\mathbf{y}_1 > \mathbf{y}_2 | \mathbf{x}) = \frac{1}{1 + \exp\left(\beta \cdot \log \frac{\pi(\mathbf{y}_2 | \mathbf{x})}{\pi_{\text{ref}}(\mathbf{y}_2 | \mathbf{x})} - \beta \cdot \log \frac{\pi(\mathbf{y}_1 | \mathbf{x})}{\pi_{\text{ref}}(\mathbf{y}_1 | \mathbf{x})}\right)}, \quad (2.17)$$

where β is a parameter controlling the sensitivity of the preference function. The inclusion of π_{ref} in DPO is to guide the fine-tuned model to follow the output of the original model.

2.4.4 Parameter-Efficient Fine-Tuning

Parameter-efficient fine-tuning (PEFT) refers to techniques to efficiently adapt a pre-trained LLM to solve new tasks with parameter efficiency. That is, training as few parameters as possible to reduce the computation and memory costs. To achieve parameter efficiency, PEFT only updates the assigned parameters while keeping the parameters of the pre-trained LLMs intact (i.e., freezing the pre-trained model weights). We list the key PEFT methods as follows.

- *Adapter* [314] introduces additional trainable layers to a transformer block of a pre-trained LLM for adaptation.
- *LoRA* (low-rank adaptation) [322] introduces a low-rank matrix AB^T to the pre-trained weight matrix W in a transformer block to update its parameters by $W \leftarrow W + AB^T$, where $W \in \mathcal{R}^{d_{\text{in}} \times d_{\text{out}}}$, $A \in \mathcal{R}^{d_{\text{in}} \times r}$, and $B \in \mathcal{R}^{d_{\text{out}} \times r}$. The value r ($r \ll \min\{d_{\text{in}}, d_{\text{out}}\}$) is the rank of the matrix AB^T . It controls the number of trainable parameters in the model updates and the memory/computation cost.
- *Hard prompt tuning* (sometimes also called *prompt engineering* or *prompt design*) refers to techniques and engineering efforts to optimize the context (e.g., the system prompt and/or the user prompt) adhering to the input format of a foundation model (e.g., textual prompts) to improve the model performance, which includes advanced methods such as in-context learning (ICL) and chain-of-thought (CoT) [888]. Note that hard prompt tuning usually does not involve any trainable parameters.
- *Soft prompt tuning* describes the category of PEFT techniques that add trainable token embeddings (i.e., soft prompts) to pre-trained LLMs.
- *Prompt tuning* [452] originates from the idea of freezing the entire pre-trained model and only allowing an additional k tunable (trainable) tokens to be prepended to the input text embeddings, as a type of soft prompt tuning methods. Let $X_e \in \mathcal{R}^{n \times h}$ denote the original input token embeddings, where n is the number of tokens and h is the size of the embedding vector. Prompt tuning in [452] adds trainable token embeddings (soft prompts) $P_e \in \mathcal{R}^{k \times h}$ such that the modified input token embeddings become $[P_e; X_e] \in \mathcal{R}^{(k+n) \times h}$. It is also worth noting that in the broader sense, the term “prompt tuning” is also used to describe general soft prompt tuning methods beyond the input token embeddings.
- *Prefix tuning* [483] adds trainable token embeddings (possibly with fully connected layers for dimension matching) to the input of each transformer block instead of only the input embeddings.

Chapter 3

Learning and Generalization of Vision Transformers



Abstract Vision Transformers (ViTs) with self-attention modules have recently achieved great empirical success in many vision tasks. Due to non-convex interactions across layers, however, the theoretical learning and generalization analysis is non-trivial. Following the framework in Li et al. (A theoretical understanding of shallow vision transformers: Learning, generalization, and sample complexity. In: International Conference on Learning Representations (2023)) and based on a data model characterizing both label-relevant and label-irrelevant tokens, this chapter provides the theoretical analysis of training a shallow ViT, i.e., one self-attention layer followed by a two-layer perceptron, for a classification task. We characterize the sample complexity to achieve a zero generalization error. The sample complexity bound is shown to be positively correlated with the inverse of the fraction of label-relevant tokens, the token noise level, and the initial model error. We also prove that a training process using stochastic gradient descent (SGD) leads to a sparse attention map, which is a formal verification of the general intuition about the success of attention. Moreover, the result indicates that a proper token sparsification can improve the test performance by removing label-irrelevant and/or noisy tokens, including spurious correlations.

3.1 Introduction

As the backbone of Transformers [845], the self-attention mechanism [37] computes the feature representation by globally modeling long-range interactions within the input. Transformers have demonstrated tremendous empirical success in numerous areas, including natural language processing [82, 390, 678, 681], recommendation system [125, 788, 1033], and reinforcement learning [116, 362, 1024]. Starting from the advent of Vision Transformer (ViT) [194], Transformer-based models [376, 537, 817, 870] gradually replace convolutional neural network (CNN) architectures and become prevalent in vision tasks. ViT also demonstrates strong robustness and generalization compared to CNN models [643, 746]. Various techniques have been developed to train ViT efficiently. Among them, token sparsification [495, 636, 695, 800, 952] removes redundant tokens (image patches) of data to

improve the computational complexity while maintaining a comparable learning performance. For example, [495, 800] prune tokens following criteria designed based on the magnitude of the attention map. Despite the remarkable empirical success, one fundamental question about training Transformers is still vastly open, which is

Under what conditions does a Transformer achieve satisfactory generalization?

Some recent works analyze Transformers theoretically from the perspective of proved Lipschitz constant of self-attention [358, 398], properties of the neural tangent kernel [318, 937] and expressive power and Turing-completeness [61, 62, 164, 176, 192, 204, 453, 500, 971] with statistical guarantees [766, 884]. [500] showed a model complexity for the function approximation of the self-attention module. Cordonnier et al. [164] provided sufficient and necessary conditions for multi-head self-attention structures to simulate convolution layers. None of these works, however, characterize the generalization performance of the learned model theoretically. Only [204] theoretically proved that a single self-attention head can represent a sparse function of the input with a sample complexity for a generalization gap between the training loss and the test loss, but no discussion is provided regarding what algorithm to train the Transformer to achieve a desirable loss.

Following [461], this chapter presents the learning and generalization analysis of training a basic shallow Vision Transformer using stochastic gradient descent (SGD). We focus on a binary classification problem on structured data, where tokens with discriminative patterns determine the label from a majority vote, while tokens with non-discriminative patterns do not affect the labels. We train a ViT containing a self-attention layer followed by a two-layer perceptron using SGD from a proper initial model. This paper explicitly characterizes the required number of training samples to achieve a desirable generalization performance, referred to as the sample complexity. The sample complexity bound is shown to be positively correlated with the inverse of the fraction of label-relevant tokens, the token noise level, and the error from the initial model, indicating a better generalization performance on data with fewer label-irrelevant patterns and less noise from a better initial model.

Generalizing the theoretical results in [461], the authors prove the property of low-rank and sparsity of one-layer transformers by characterizing the trained model after convergence using stochastic gradient descent [467]. In [466], the authors also extend the theoretical analysis to graph transformer networks and demonstrate the roles of self-attention and positional encoding for learning and generalization.

3.2 Background and Related Work

Efficient ViT Learning To alleviate the memory and computation burden in training [194, 817, 877], various acceleration techniques have been developed other than token sparsification. Zhu et al. [1045] identifies the importance of different dimensions in each layer of ViTs and then executes model pruning. Liu et al.

[542], Lin et al. [504], and Li et al. [491] quantize weights and inputs to compress the learning model. Li et al. [457] studies automated progressive learning that automatically increases the model capacity on-the-fly. Moreover, modifications of attention modules, such as the network architecture based on local attention [153, 537, 870], can simplify the computation of global attention for acceleration.

Theoretical Analysis of Learning and Generalization of Neural Networks One line of research [234, 469, 995, 997, 1028, 1029] analyzes the generalization performance when the number of neurons is smaller than the number of training samples. The neural-tangent-kernel (NTK) analysis [19, 20, 33, 88, 136, 195, 355, 462, 1054] considers strongly overparameterized networks and eliminates the nonconvex interactions across layers by linearizing the neural network around the initialization. The generalization performance is independent of the feature distribution and cannot explain the advantages of self-attention modules.

Neural Network Learning on Structured Data Li and Liang [489] provide the generalization analysis of a fully-connected neural network when the data comes from separated distributions. Daniely and Malach [173], Shi et al. [755], Karp et al. [386], Brutzkus and Globerson [84], and Zhang et al. [994] study fully connected networks and convolutional neural networks assuming that data contains discriminative patterns and background patterns. Allen-Zhu and Li [17] illustrates the robustness of adversarial training by introducing the feature purification mechanism, in which neural networks with non-linear activation functions can memorize the data-dependent features. Shi et al. [898] extends this framework to the area of self-supervised contrastive learning. All these works consider one-hidden-layer neural networks without self-attention.

Notations Vectors are in bold lowercase, and matrices and tensors are in bold uppercase. Scalars are in normal fonts. Sets are in calligraphy font. For instance, $\mathbf{Z}$ is a matrix, and $\mathbf{z}$ is a vector. z_i denotes the i -th entry of $\mathbf{z}$, and $Z_{i,j}$ denotes the (i, j) -th entry of $\mathbf{Z}$. $[K]$ ($K > 0$) denotes the set including integers from 1 to K . We follow the convention that $f(x) = O(g(x))$ (or $\Omega(g(x))$, $\Theta(g(x))$) means that $f(x)$ increases at most, at least, or in the order of $g(x)$, respectively.

3.2.1 Problem Formulation and Learning Algorithm

We study a binary classification problem. following the common setup in [194, 376, 817]. Given N training samples $\{(X^n, y^n)\}_{n=1}^N$ generated from an unknown distribution $\mathcal{D}$ and a fair initial model, the goal is to find an improved model that maps X to y for any $(X, y) \sim \mathcal{D}$. Here each data point contains L tokens $\mathbf{x}_1^n, \mathbf{x}_2^n, \dots, \mathbf{x}_L^n$, i.e., $\mathbf{X}^n = [\mathbf{x}_1^n, \dots, \mathbf{x}_L^n] \in \mathbb{R}^{d \times L}$, where each token is d -dimensional and unit-norm. $y^n \in \{+1, -1\}$ is a scalar. A token can be an image patch [194]. We consider a general setup that also applies to token sparsification, where some tokens are set to zero to reduce the computational time. Let $\mathcal{S}^n \subseteq [L]$

denote the set of indices of remaining tokens in X^n after sparsification. Then $|\mathcal{S}^n| \leq L$, and $\mathcal{S}^n = [L]$ without token sparsification.

Learning is performed over a basic shallow Vision Transformer, a neural network with a single-head self-attention layer and a two-layer fully connected network, as shown in (3.1). This is a simplified model of practical Vision Transformers [194] to avoid unnecessary complications in analyzing the most critical component of ViTs, the self-attention.

$$F(X^n) = \frac{1}{|\mathcal{S}^n|} \sum_{l \in \mathcal{S}^n} \mathbf{a}_{(l)}^\top \text{Relu}(\mathbf{W}_O \mathbf{W}_V X^n \text{softmax}(X^n^\top \mathbf{W}_K^\top \mathbf{W}_Q x_l^n)), \quad (3.1)$$

where the queue weights $\mathbf{W}_Q$ in $\mathbb{R}^{m_b \times d}$, the key weights $\mathbf{W}_K$ in $\mathbb{R}^{m_b \times d}$, and the value weights $\mathbf{W}_V$ in $\mathbb{R}^{m_a \times d}$ in the attention unit are multiplied with X^n to obtain the queue vector $\mathbf{W}_Q X^n$, the key vector $\mathbf{W}_K X^n$, and the value vector $\mathbf{W}_V X^n$, respectively [845]. $\mathbf{W}_O$ is in $\mathbb{R}^{m \times m_a}$ and $\mathbf{A} = (\mathbf{a}_{(1)}, \mathbf{a}_{(2)}, \dots, \mathbf{a}_{(L)})$ where $\mathbf{a}_{(l)} \in \mathbb{R}^m$, $l \in [L]$ are the hidden-layer and output-layer weights of the two-layer perceptron, respectively. m is the number of neurons in the hidden layer. $\text{Relu} : \mathbb{R}^m \rightarrow \mathbb{R}^m$ where $\text{Relu}(\mathbf{x}) = \max\{\mathbf{x}, 0\}$. $\text{softmax} : \mathbb{R}^L \rightarrow \mathbb{R}^L$ where $\text{softmax}(\mathbf{x}) = (e^{x_1}, e^{x_2}, \dots, e^{x_L}) / \sum_{i=1}^L e^{x_i}$. Let $\psi = (\mathbf{A}, \mathbf{W}_O, \mathbf{W}_V, \mathbf{W}_K, \mathbf{W}_Q)$ denote the set of parameters to train. The training problem minimizes the empirical risk $f_N(\psi)$,

$$\min_{\psi} : f_N(\psi) = \frac{1}{N} \sum_{n=1}^N \ell(X^n, y^n; \psi), \quad (3.2)$$

where $\ell(X^n, y^n; \psi)$ is the Hinge loss function, i.e.,

$$\ell(X^n, y^n; \psi) = \max\{1 - y^n \cdot F(X^n), 0\}. \quad (3.3)$$

The generalization performance of a learned model ψ is evaluated by the population risk $f(\psi)$, where

$$f(\psi) = f(\mathbf{A}, \mathbf{W}_O, \mathbf{W}_V, \mathbf{W}_K, \mathbf{W}_Q) = \mathbb{E}_{(X, y) \sim \mathcal{D}}[\max\{1 - y \cdot F(X), 0\}]. \quad (3.4)$$

The training problem (3.2) is solved via a mini-batch stochastic gradient descent (SGD), as summarized in Algorithm 1. At iteration t , $t = 0, 1, 2, \dots, T - 1$, the gradient is computed using a mini-batch $\mathcal{B}_t$ with $|\mathcal{B}_t| = B$. The step size is η .

Algorithm 1 Training with SGD

-
- 1: **Input:** Training data $\{(X^n, y^n)\}_{n=1}^N$, the step size η , the total number of iterations T , batch size B .
 - 2: **Initialization:** Every entry of $\mathbf{W}_O^{(0)}$ from $\mathcal{N}(0, \xi^2)$, and every entry of $\mathbf{a}_{(l)}^{(0)}$ from $\text{Uniform}([+\frac{1}{\sqrt{m}}, -\frac{1}{\sqrt{m}}])$. $\mathbf{W}_V^{(0)}$, $\mathbf{W}_K^{(0)}$ and $\mathbf{W}_Q^{(0)}$ from a pre-trained model.
 - 3: **Stochastic Gradient Descent:** for $t = 0, 1, \dots, T - 1$ and $\mathbf{W}^{(t)} \in \{\mathbf{W}_O^{(t)}, \mathbf{W}_V^{(t)}, \mathbf{W}_K^{(t)}, \mathbf{W}_Q^{(t)}\}$

$$\mathbf{W}^{(t+1)} = \mathbf{W}^{(t)} - \eta \cdot \frac{1}{B} \sum_{n \in \mathcal{B}_t} \nabla_{\mathbf{W}^{(t)}} \ell(X^n, y^n; \mathbf{W}_O^{(t)}, \mathbf{W}_V^{(t)}, \mathbf{W}_K^{(t)}, \mathbf{W}_Q^{(t)}) \quad (3.5)$$

- 4: **Output:** $\mathbf{W}_O^{(T)}, \mathbf{W}_V^{(T)}, \mathbf{W}_K^{(T)}, \mathbf{W}_Q^{(T)}$.
-

3.3 Theoretical Characterization of Transformers

3.3.1 Main Theoretical Insights

Before formally introducing the data model and main theory in [461], we first summarize the major insights. We consider a data model where tokens are noisy versions of *label-relevant* patterns that determine the data label and *label-irrelevant* patterns that do not affect the label. α_* is the fraction of label-relevant tokens. σ represents the initial model error, and τ characterizes the token noise level.

(P1). A Convergence and Sample Complexity Analysis of SGD to Achieve Zero Generalization Error We prove SGD with a proper initialization converges to a model with zero generalization error. The required number of iterations is proportional to $1/\alpha_*$ and $1/(\Theta(1) - \sigma - \tau)$. Our sample complexity bound is linear in α_*^{-2} and $(\Theta(1) - \sigma - \tau)^{-2}$. Therefore, the learning performance is improved, in the sense of a faster convergence and fewer training samples to achieve a desirable generalization, with a larger fraction of label-relevant patterns, a better initial model, and less token noise.

(P2). A Theoretical Characterization of Increased Sparsity of the Self-Attention Module During Training We prove that the attention weights, which are softmax values of each token in the self-attention module, become increasingly sparse during the training, with non-zero weights concentrated at label-relevant tokens. This formally justifies the general intuition that the attention layer makes the neural network focus on the most important part of data.

(P3). A Theoretical Guideline of Designing Token Sparsification Methods to Reduce Sample Complexity Our sample complexity bound indicates that the required number of samples to achieve zero generalization can be reduced if a token sparsification method removes some label-irrelevant tokens (reducing α_*), or tokens with large noise (reducing σ), or both. This insight provides a guideline to design proper token sparsification methods.

(P4). A New Theoretical Framework to Analyze the Nonconvex Interactions in Shallow ViTs This paper develops a new framework to analyze ViTs based on a more general data model than existing works like [84, 386, 898]. Compared with the nonconvex interactions in three-layer feedforward neural networks, analyzing ViTs has technical challenges that the softmax activation is highly non-linear, and the gradient computation on token correlations is complicated. We develop new tools to handle this problem by exploiting structures in the data and proving that SGD iterations increase the magnitude of label-relevant tokens only rather than label-irrelevant tokens. This theoretical framework is of independent interest and can be potentially applied to analyze different variants of Transformers and attention mechanisms.

3.3.2 Data Model

We assume a data model such that there are M ($\Theta(1) < M < m_a, m_b = \Theta(M)$) distinct patterns $\{\mu_1, \mu_2, \dots, \mu_M\}$ in $\mathbb{R}^d$, where μ_1, μ_2 are *discriminative patterns* that determine the binary labels, and the remaining $M - 2$ patterns $\mu_3, \mu_4, \dots, \mu_M$ are *non-discriminative patterns* that do not affect the labels. Let $\kappa = \min_{1 \leq i \neq j \leq M} \|\mu_i - \mu_j\| > 0$ denote the minimum distance between patterns. Each token x_l^n of X^n is a noisy version of one of the patterns, i.e.,

$$\min_{j \in [M]} \|x_l^n - \mu_j\| \leq \tau, \quad (3.6)$$

and the noise level $\tau < \kappa/4$. We take $\kappa - 4\tau$ as $\Theta(1)$ for the simplicity of presentation.

The label y^n is determined by the tokens that correspond to discriminative patterns through a majority vote. If the number of tokens that are noisy versions of μ_1 is larger than the number of tokens that correspond to μ_2 in X^n , then $y^n = 1$. In this case that the label $y^n = 1$, the tokens that are noisy μ_1 are referred to as *label-relevant* tokens, and the tokens that are noisy μ_2 are referred to as *confusion* tokens. Similarly, if there are more tokens that are noisy μ_2 than those that are noisy μ_1 , the former are label-relevant tokens, the latter are confusion tokens, and $y^n = -1$. All other tokens that are not label-relevant are called label-irrelevant tokens.

Let α_* and $\alpha_\#$ as the average fraction of the label-relevant and the confusion tokens over the distribution $\mathcal{D}$, respectively. We consider a balanced dataset. Let $\mathcal{D}_+ = \{(X^n, y^n) | y^n = +1, n \in [N]\}$ and $\mathcal{D}_- = \{(X^n, y^n) | y^n = -1, n \in [N]\}$ denote the sets of positive and negative labels, respectively. Then $|\mathcal{D}_+| - |\mathcal{D}_-| = O(\sqrt{N})$.

Our model is motivated by and generalized from those used in the state-of-art analysis of neural networks on structured data [84, 386, 489]. All the existing models require that only one discriminative pattern exists in each sample, i.e., either μ_1 or μ_2 , but not both, while our model allows both patterns to appear in the same sample.

Table 3.1 Some important notations

σ	Initialization error for value vectors	δ	Initialization error for query and key vectors
κ	Minimum of $\ \mu_i - \mu_j\ $ for any $i, j \in [M], i \neq j$.	τ	Token noise level
M	Total number of patterns	m	The number of neurons in W_O
α_*	Average fraction of label-relevant tokens	$\alpha_\#$	Average fraction of confusion tokens

3.3.3 Formal Theoretical Results

Before presenting our main theory below, we first characterize the behavior of the initial model through Assumption 3.1. Some important notations are summarized in Table 3.1.

Assumption 3.1 Assume $\max(\|W_V^{(0)}\|, \|W_K^{(0)}\|, \|W_Q^{(0)}\|) \leq 1$ without loss of generality. There exist three (not necessarily different) sets of orthonormal bases $\mathcal{P} = \{p_1, p_2, \dots, p_M\}$, $\mathcal{Q} = \{q_1, q_2, \dots, q_M\}$, and $\mathcal{R} = \{r_1, r_2, \dots, r_M\}$, where $p_l \in \mathbb{R}^{m_a}$, $q_l, r_l \in \mathbb{R}^{m_b}$, $\forall l \in [M]$, $q_1 = r_1$, and $q_2 = r_2$ ¹ such that

$$\|W_V^{(0)} \mu_j - p_j\| \leq \sigma, \quad \|W_K^{(0)} \mu_j - q_j\| \leq \delta, \quad \text{and} \quad \|W_Q^{(0)} \mu_j - r_j\| \leq \delta. \quad (3.7)$$

hold for some $\sigma = O(1/M)$ and $\delta < 1/2$.

Assumption 3.1 characterizes the distance of query, key, and value vectors of patterns $\{\mu_j\}_{j=1}^M$ to orthonormal vectors. The requirement on δ is minor because δ can be in the same order as $\|\mu_j\|$.

Theorem 3.1 (Generalization of ViT [461]) Suppose Assumption 3.1 holds; $\tau \leq \min(\sigma, \delta)$; a sufficiently large model with

$$m \gtrsim M^2 \log N, \quad (3.8)$$

the average fraction of label-relevant patterns satisfies

$$\alpha_* \geq \frac{\alpha_\#}{e^{-(\delta+\tau)}(1 - (\sigma + \tau))}, \quad (3.9)$$

and the mini-batch size and the number of sampled tokens of each data X^n , $n \in [N]$ satisfy

$$B \geq \Omega(1), \quad |S^n| \geq \Omega(1) \quad (3.10)$$

¹ The condition $q_1 = r_1$ and $q_2 = r_2$ is to eliminate the trivial case that the initial attention value is very small. This condition can be relaxed but we keep this form to simplify the representation.

Then, after T number of iterations such that

$$T = \Theta(\eta^{-3/5} \alpha_*^{-1}), \quad (3.11)$$

as long as the number of training samples N satisfies

$$N \geq \Omega\left(\frac{1}{(\alpha_* - c'(1 - \zeta) - c''(\sigma + \tau))^2}\right) \quad (3.12)$$

for some constant $c', c'' > 0$, and $\zeta \gtrsim 1 - \eta^{10}$, with a probability of at least 0.99, the returned model achieves zero generalization error as

$$f(\mathbf{A}^{(0)}, \mathbf{W}_O^{(T)}, \mathbf{W}_V^{(T)}, \mathbf{W}_K^{(T)}, \mathbf{W}_Q^{(T)}) = 0 \quad (3.13)$$

Theorem 3.1 characterizes under what condition of the data the neural network with self-attention in (3.1) trained with Algorithm 1 can achieve zero generalization error. To show that the self-attention layer can improve the generalization performance by reducing the required sample complexity to achieve zero generalization error, we also quantify the sample complexity when there is no self-attention layer in the following proposition.

Proposition 3.1 (Generalization Without Self-Attention [461]) Suppose assumptions in Theorem 3.1 hold. When there is no self-attention layer, i.e., $\mathbf{W}_K$ and $\mathbf{W}_Q$ are not updated during the training, if N satisfies

$$N \geq \Omega\left(\frac{1}{(\alpha_*(\alpha_* - \sigma - \tau))^2}\right) \quad (3.14)$$

then after T iterations with T in (3.11), the returned model achieves zero generalization error as

$$f(\mathbf{A}^{(0)}, \mathbf{W}_O^{(T)}, \mathbf{W}_V^{(T)}, \mathbf{W}_K^{(0)}, \mathbf{W}_Q^{(0)}) = 0 \quad (3.15)$$

Remark 3.1 (Advantage of the Self-Attention Layer) Because $m \gg m_a, m_b, d$, the number of trainable parameter remains almost the same with or without updating the attention layer. Combining Theorem 3.1 and Proposition 3.1, we can see that with the additional self-attention layer, the sample complexity² is reduced by a factor $1/\alpha_*^2$ with an approximately equal number of network parameters.

² The sample complexity bounds in (3.12) and (3.14) are sufficient but not necessary. Thus, rigorously speaking, one can not compare two cases based on sufficient conditions only. In our analysis, however, these two bounds are derived with exactly the same technique with the only difference in handling the self-attention layer. Therefore, we believe it is fair to compare these two bounds to show the advantage of ViT.

Remark 3.2 (Generalization Improvement by Token Sparsification) (3.12) and (3.11) show that the sample complexity N and the required number of iterations T scale with $1/\alpha_*^2$ and $1/\alpha_*$, respectively. Then, increasing α_* , the fraction of label-relevant tokens, can reduce the sample complexity and speed up the convergence. Similarly, N and T scale with $1/(\Theta(1) - \tau)^2$ and $1/(\Theta(1) - \tau)$. Then decreasing τ , the noise in the tokens, can also improve the generalization. Note that a properly designed token sparsification method can both increase α_* by removing label-irrelevant tokens and decrease τ by removing noisy tokens, thus improving the generalization performance.

Remark 3.3 (Impact of the Initial Model) The initial model $\mathbf{W}_V^{(0)}$, $\mathbf{W}_K^{(0)}$, $\mathbf{W}_Q^{(0)}$ affects the learning performance through σ and δ , which decrease as the initial model is improved. Then from (3.12) and (3.11), the sample complexity reduces and the convergence speeds up for a better initial model.

Proposition 3.2 shows that the attention weights are increasingly concentrated on label-relevant tokens during the training. Proposition 3.2 is a critical component in proving Theorem 3.1 and is of independent interest.

Proposition 3.2 (Concentration of Attention Weights [461]) *The attention weights for each token become increasingly concentrated on those correlated with tokens of the label-relevant pattern during the training, i.e.,*

$$\begin{aligned} \sum_{i \in \mathcal{S}_*^n} \text{softmax}(\mathbf{X}^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{x}_l^n)_i &= \sum_{i \in \mathcal{S}_*^n} \frac{\exp(\mathbf{x}_i^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{x}_l^n)}{\sum_{r \in \mathcal{S}^n} \exp(\mathbf{x}_r^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{x}_l^n)} \\ &\rightarrow \begin{cases} 1 - \eta^C, & \text{if } \mathbf{x}_l^n \text{ corresponds to either of } \boldsymbol{\mu}_1, \boldsymbol{\mu}_2 \\ 1 - \eta^C - e^{-(\delta+\tau)} & \text{if } \mathbf{x}_l^n \text{ corresponds to one of } \boldsymbol{\mu}_3, \dots, \boldsymbol{\mu}_M \end{cases} \end{aligned} \quad (3.16)$$

at a sublinear rate of $O(1/t)$ when t is large for a large $C > 0$ and all $l \in \mathcal{S}^n$ and $n \in [N]$.

3.4 Performance Evaluation

Experiment Setup We verify the theoretical bounds in Theorem 3.1 on synthetic data. We set the dimension of data and attention embeddings to be $d = m_a = m_b = 10$. Let $c_0 = 0.01$. Let the total number of patterns $M = 5$, and $\{\boldsymbol{\mu}_1, \boldsymbol{\mu}_2, \dots, \boldsymbol{\mu}_M\}$ be a set of orthonormal bases. To satisfy Assumption 3.1, we generate every token that is a noisy version of $\boldsymbol{\mu}_i$ from a Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}_i, c_0^2 \cdot \mathbf{I})$ with the mean $\boldsymbol{\mu}_i$ and covariance $c_0^2 \mathbf{I}$, where $\mathbf{I} \in \mathbb{R}^d$ is the identity matrix. We set $\mathbf{W}_Q^{(0)} = \mathbf{W}_Q^{(0)}$ $= \delta^2 \mathbf{I} / c_0^2$, $\mathbf{W}_V^{(0)} = \sigma^2 \mathbf{U} / c_0^2$, and each entry of $\mathbf{W}_O^{(0)}$ follows $\mathcal{N}(0, \xi^2)$, where $\mathbf{U}$ is an $m_a \times m_a$ orthonormal matrix, and $\xi = 0.01$. The number of neurons m of

W_O is 1000. We set the ratio of different patterns the same among all the data for simplicity.

Numerical Results We summarize the key findings as follows.

Sample Complexity and Convergence Rate We first study the impact of the fraction of the label-relevant patterns α_* on the sample complexity. Let the number of tokens after sparsification be $|\mathcal{S}^n| = 100$, the initialization error $\sigma = 0.1$ and $\delta = 0.2$. The fraction of non-discriminative patterns is fixed to be 0.5. We implement 20 independent experiments with the same α_* and N and record the Hinge loss values of the testing data. An experiment is successful if the testing loss is smaller than 10^{-3} . Figure 3.1a shows the success rate of these experiments. A black block means that all the trials fail. A white block means that they all succeed. The sample complexity is indeed almost linear in α_*^{-2} , as predicted in (3.12). We next explore the impact on σ . Set $\alpha_* = 0.3$ and $\alpha_{\#} = 0.2$. The number of tokens after sparsification is fixed at 50 for all the data. Figure 3.1b shows that $1/\sqrt{N}$ is linear in $\Theta(1) - \sigma$, matching our theoretical prediction in (3.12). The result on the noise level τ is similar to Fig. 3.1b, and we skip it here. In Fig. 3.2, we verify the number of iterations T against α_*^{-1} in (3.11) where we set $\sigma = 0.1$ and $\delta = 0.4$.

Advantage of Self-Attention To verify Proposition 3.1, we compare the performance on ViT in (3.1) and on the same network with W_K and W_Q fixed during the training, i.e., a three-layer CNN. Compared with ViT, the number of trainable parameters in CNN is reduced by only 1%. Figure 3.3 shows the sample complexity of CNN is almost linear in α_*^{-4} as predicted in (3.14). Compared with Fig. 3.2, the sample complexity significantly increases for small α_* , indicating a much worse generalization of CNN.

Attention Map We then evaluate the evolution of the attention map during the training. Let $|\mathcal{S}^n| = 50$ for all $n \in [N]$. The number of training samples is $N = 200$. $\sigma = 0.1$, $\delta = 0.2$, $\alpha_* = 0.5$, $\alpha_{\#} = 0.05$. In Fig. 3.4, the red line with asterisks shows that the sum of attention weights on label-relevant tokens, i.e., the left side of (3.16) averaged over all l , indeed increases to be close to 1 when the number of iterations increases. Correspondingly, the sum of attention weights on other tokens decreases




Fig. 3.1 The impact of α_* and σ on sample complexity





Fig. 3.2 The number of iterations against α_*^{-1}

Fig. 3.3 Comparison of ViT and CNN



to be close to 0, as shown in the blue line with squares. This verifies Lemma 3.2 on a sparse attention map.

Token Sparsification We verify the improvement by token sparsification in Fig. 3.5. The experiment is duplicated 20 times. The number of training samples $N = 80$. Let $|\mathcal{S}^n| = 50$ for all $n \in [N]$. Set $\sigma = 0.1$, $\delta = 0.5$, $\alpha_* = 0.6$, $\alpha_{\#} = 0.05$. If we apply random sampling over all tokens, the performance cannot be improved as shown in the red curve because α_* and σ do not change. If we remove either label-irrelevant tokens or tokens with significant noise, the testing loss decreases, as indicated in the blue and black curves. This justifies our insight **P3** on token sparsification.

Fig. 3.4 Concentration of attention weights





Fig. 3.5 Impact of token sparsification on testing loss



Chapter 4

Formalizing In-Context Learning in Transformers



Abstract Transformer-based large language models have shown impressive in-context learning (ICL) capabilities, where a pre-trained model can handle new tasks without fine-tuning by simply augmenting the query with some input-output examples from that task. Despite the empirical success, the mechanics of how to train a Transformer to achieve ICL and the corresponding ICL capacity are not straightforward, due to the technical challenges of analyzing the nonconvex training problems resulting from the nonlinear self-attention and nonlinear activation in transformers. This chapter provides the theoretical analysis of the training dynamics of Transformers with nonlinear self-attention and nonlinear MLP, together with the ICL generalization capability of the resulting model, as proved in Li et al. (How do nonlinear transformers learn and generalize in in-context learning? In: International Conference on Machine Learning (2024)). Focusing on a group of binary classification tasks, we train transformers using data from a subset of these tasks and quantify the impact of various factors on the ICL generalization performance on the remaining unseen tasks with and without data distribution shifts. We also analyze how different components in the learned Transformers contribute to the ICL performance. Furthermore, we provide the theoretical analysis of how model pruning affects ICL performance and prove that proper magnitude-based pruning can have a minimal impact on ICL while reducing inference costs.

4.1 Introduction

Transformers serve as the backbone architecture for a wide range of modern, large-scale foundation models, including prominent language models like GPT-3 [82], PaLM [148], LLaMa [818], as well as versatile visual and multi-modal models such as CLIP [675], DALL-E [689], and GPT-4 [619]. One intriguing capability exhibited by certain large language models (LLMs) is known as “**in-context learning**” (ICL) [82]. Given a pre-trained model $F(\Psi)$, parameterized by weights Ψ , the conventional approach fine-tunes Ψ separately for each downstream task using data from that task. In contrast, ICL allows $F(\Psi)$ to handle multiple unseen tasks simultaneously without any fine-tuning. The work in [248] is the first paper to

mathematically formulate ICL. Briefly speaking, to predict $f(\mathbf{x}_{\text{query}})$ of a query input $\mathbf{x}_{\text{query}}$ for a new task represented by the label function f , ICL augments $\mathbf{x}_{\text{query}}$ by l example input-output pairs $(\mathbf{x}_i, f(\mathbf{x}_i))_{i=1}^l$. The resulting so-called *prompt* is sent to the model $F(\Psi)$, and, surprisingly, the model can output a prediction close to $f(\mathbf{x}_{\text{query}})$. Thus, ICL is an efficient alternative to the resource-consuming fine-tuning methods. ICL has shown outstanding performance in multiple tasks in practice, including question answering [518, 915], natural language inference [510, 915], text generation [82, 550], etc.

In parallel, model pruning [282, 896] can reduce the inference cost by removing some weights after training. It has been extensively evaluated in various applications. Among various pruning techniques, such as gradient methods [594] and reconstruction error minimization [552], magnitude-based pruning [896] is the most popular approach due to its simplicity and demonstrated promising empirical results. A few recent works [232, 541, 558, 792] also explore the pruning of LLMs to preserve their ICL capacity while accelerating the inference. Despite the empirical success of ICL, one fundamental and theoretical question is less investigated, which is:

How can a Transformer be trained to perform ICL and generalize in and out of domain successfully and efficiently?

Some recent works attempt to answer this question for linear regression tasks [486, 988]. Specifically, [486] investigate the generalization gap and stability of ICL. Zhang et al. [988] explore the training and generalization of ICL with Transformers, especially with distribution shifts during inference. Wu et al. [909] studies the required number of pre-training tasks for a desirable ICL property. Huang et al. [341] characterizes the training dynamics using Transformers with softmax attention and linear MLP. However, these results are either built upon simplified Transformer models by ignoring nonlinear self-attention [909, 988] or nonlinear activation in the multilayer perceptron (MLP) [341, 909, 988] or cannot characterize how to train a model to achieve the desirable ICL capability with distribution-shifted data [341, 486, 909].

In [464], the authors provide a theoretical analysis of the training dynamics of Transformers with nonlinear self-attention and nonlinear MLP, together with the ICL generalization capability of the resulting model. Moreover, they also provide a theoretical analysis of the impact of model pruning on ICL performance. Focusing on a group of binary classification tasks, we show that training a Transformer using prompts from a subset of these tasks can return a model with the ICL capability to generalize to the rest of these tasks.

As will be presented throughout this chapter, the key insights from [464] include:

1. **A theoretical characterization of how to train Transformers to enhance their ICL capability.** By considering a data model where input data include both relevant patterns that determine the labels and irrelevant patterns that do not affect the labels, [464] quantifies how the training and the resulting ICL generalization performance are affected by various factors, such as the magnitude

of relevant features and the fraction of context examples that contain the same relevant pattern as the new query. In addition to proving the ICL capability of the learned Transformer to generalize to new binary tasks based on the relevant patterns that appear in the training data, it is also proved in [464] that the ICL capability to generalize to tasks based on patterns that are linear combinations of the relevant patterns and are unseen in the training data.

2. **Expand the theoretical understanding of the mechanism of the ICL capability of Transformers.** Li et al. [464] proves that when sending a prompt to a properly trained Transformer, the attention weights are concentrated on contexts that share the same relevant pattern as the query. Then, the ReLU MLP layer promotes the label embedding of these examples, thus making the correct prediction for the query. Similar insights have appeared in [341]. Li et al. [464] expands the analysis to Transformers with nonlinear MLP layers and new tasks with a data distribution shift.
3. **Theoretical justification of magnitude-based pruning in preserving ICL.** Based on the characterization of the trained Transformer, [464] also provides a theoretical analysis of the ICL inference performance when the trained model is pruned by removing neurons in the MLP layer. It can be shown that pruning a set of neurons with a small magnitude has little effect on the generalization while pruning the remaining neurons leads to a large generalization error growing with the pruning rate.

Delving into the chain-of-thoughts (CoT) reasoning capabilities of LLMs [889], the authors in [463] extend the theoretical analysis to study how transformers acquire CoT ability.

4.2 Background and Related Work

Expressive Power of ICL Some existing works study the expressive power of Transformers to implement algorithms via ICL. Akyürek et al. [16] and Von Oswald et al. [849] demonstrate that Transformers conduct gradient descent during the forward pass of Transformers with prompts as inputs. Ahn et al. [11] and Cheng et al. [139] extend the conclusion to preconditioned and functional gradient descent via ICL. Garg et al. [248], Bai et al. [40], and Guo et al. [278] show the existence of Transformers that can implement a broad class of machine learning algorithms in context.

The Optimization and Generalization of Transformers Beyond in-context learning, there are several other works about the optimization and generalization analysis of fine-tuning or prompt tuning on Transformers. Jelassi et al. [364], Li et al. [461, 465], and Luo et al. [553] study the generalization of one-layer Transformer by assuming spatial association or the majority voting of tokens. Li et al. [466] investigate the effect of the relative positional encoding in training for Graph Transformer. Li et al. [487] delve into how one-layer Transformers

learn semantic structure. Oymak et al. [630] depict the trajectory of prompt tuning of attention networks. Tarzanagh et al. [805, 806] characterize that the gradient updates of the prompt or weights converge to a max-margin SVM solution. Tian et al. [812, 813] probe the training dynamics of Transformers for the next token prediction problem given infinitely long sequences.

Theoretical Generalization Analysis of Pruning A few recent works consider analyzing the generalizations performance of model pruning theoretically. For example, [996] study the sample complexity of training a pruned network with a given sparse ground truth weight. Yang and Wang [939] investigate the neural tangent kernel of the pruned model. Zhang et al. [994] and Yang et al. [938] consider the generalization using magnitude pruning under a feature learning framework. However, these works are built on convolutional neural networks, and no theoretical works are for LLM or Transformer-based models.

4.2.1 Formalizing In-Context Learning with Transformers

We study the optimization and generalization of binary classification problems for in-context learning. Consider a query $\mathbf{x}_{query}$ and its label z . Define a set of binary classification tasks $\mathcal{T}$, consisting of multiple task functions. The label $z \in \{+1, -1\}$ is mapped from $\mathbf{x}_{query} \in \mathbb{R}^{d_X}$ through a task f that is randomly chosen from $\mathcal{T}$, i.e., $z = f(\mathbf{x}_{query}) \in \{+1, -1\}$, $f \in \mathcal{T}$.

Training to Enhance ICL Capability Following the framework of training for ICL in [16, 40, 248], we consider the problem of training such that the model has the ICL capability to generalize to new tasks using prompts. The idea is to update the model during the training process using pairs of the constructed prompt, embedded as $\mathbf{P}$ for the query $\mathbf{x}_{query}$, and its label $f(\mathbf{x}_{query})$. We start by formulating $\mathbf{P}$ and then introduce the learning model in this section.

Following [341, 849, 988], the prompt embedding $\mathbf{P}$ of query $\mathbf{x}_{query}$ is formulated as:

$$\mathbf{P} = \begin{pmatrix} \mathbf{x}_1 & \mathbf{x}_2 & \cdots & \mathbf{x}_l & \mathbf{x}_{query} \\ \mathbf{y}_1 & \mathbf{y}_2 & \cdots & \mathbf{y}_l & \mathbf{0} \end{pmatrix} \quad (4.1)$$

$$:= (\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_{query}) \in \mathbb{R}^{(d_X + d_Y) \times (l+1)},$$

where the last column of $\mathbf{P}$, denoted by $\mathbf{p}_{query}$, includes the query $\mathbf{x}_{query}$ with padding zeros, and the first l columns are the contexts for $\mathbf{x}_{query}$. We respectively call $\mathbf{x}_i$ and $\mathbf{y}_i, i \in [l]$ context inputs and outputs, where l is also known as the prompt length. Let $\text{Embd}(\cdot)$ be the embedding function of each context output. $\mathbf{y}_i \in \mathbb{R}^{d_Y}$ in (4.1) is defined as $\mathbf{y}_i = \text{Embd}(f(\mathbf{x}_i))$. Hence, $\mathbf{P}$ is a function of f . The first d_X dimensions of $\mathbf{p}_i$ are referred to as the feature embedding, while the last d_Y dimensions are called the label embedding.

The **learning model** is a single-head, one-layer Transformer with one self-attention layer and one two-layer perceptron. Mathematically, it can be written as

$$F(\Psi; \mathbf{P}) = \mathbf{a}^\top \text{Relu}(\mathbf{W}_O \sum_{i=1}^l \mathbf{W}_V \mathbf{p}_i \cdot \text{attn}(\Psi; \mathbf{P}, i)), \quad (4.2)$$

$$\text{attn}(\Psi; \mathbf{P}, i) = \text{softmax}((\mathbf{W}_K \mathbf{p}_i)^\top \mathbf{W}_Q \mathbf{p}_{\text{query}}),$$

where $\mathbf{W}_Q, \mathbf{W}_K \in \mathbb{R}^{m_a \times (d_X + d_Y)}$, $\mathbf{W}_V \in \mathbb{R}^{m_b \times (d_X + d_Y)}$ are the embedding matrices for queries, keys, and values, respectively, and $\mathbf{W}_O \in \mathbb{R}^{m \times m_b}$ and $\mathbf{a} \in \mathbb{R}^m$ are parameters in the MLP layer. Here, $\text{softmax}((\mathbf{W}_K \mathbf{p}_i)^\top \mathbf{W}_Q \mathbf{p}_{\text{query}}) = e^{(\mathbf{W}_K \mathbf{p}_i)^\top \mathbf{W}_Q \mathbf{p}_{\text{query}}} / \sum_{j=1}^l e^{(\mathbf{W}_K \mathbf{p}_j)^\top \mathbf{W}_Q \mathbf{p}_{\text{query}}}$. $\Psi := \{\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V, \mathbf{W}_O, \mathbf{a}\}$ denotes the set of all model weights. Typically, $\min(m_a, m_b) > d_X + d_Y$.

The **training problem** to enhance the ICL capability solves the empirical risk minimization problem,

$$\min_{\Psi} R_N(\Psi) := \frac{1}{N} \sum_{n=1}^N \ell(\Psi; \mathbf{P}^n, z^n), \quad (4.3)$$

using N pairs of prompt embedding and label pairs $\{\mathbf{P}^n, z^n\}_{n=1}^N$. For the n -th pair, $\mathbf{x}_{\text{query}}^n$ and the context input $\mathbf{x}_i^n$ are all sampled from an unknown distribution $\mathcal{D}$, the task f^n is sampled from $\mathcal{T}$, and $\mathbf{P}^n$ is constructed following (4.1). The loss function is a Hinge loss, i.e., $\ell(\Psi; \mathbf{P}^n, z^n) = \max\{0, 1 - z^n \cdot F(\Psi; \mathbf{P}^n)\}$, where $F(\Psi; \mathbf{P}^n)$ is defined in (4.2). Let $\mathcal{T}_{tr} = \bigcup_{n=1}^N f^n$ denote the set of tasks that appear in the training samples. Note that $\mathcal{T}_{tr} \subset \mathcal{T}$, and (4.3) is a *multi-task learning* problem when $|\mathcal{T}_{tr}| > 1$.

Generalization Evaluation We define two quantities to evaluate the ICL generalization performance to new tasks as follows.

In-Domain Generalization If the testing queries are also drawn from $\mathcal{D}$ and all the testing tasks are drawn from $\mathcal{T}$, we call it *in-domain* inference, and the in-domain generalization error is defined as¹

$$\mathbb{E}_{\mathbf{x}_{\text{query}} \sim \mathcal{D}, f \in \mathcal{T}} [\ell(\Psi; \mathbf{P}, z)], \quad (4.4)$$

¹ In terms of evaluating generalization on unseen tasks, (4.4) is almost equivalent to replacing $f \in \mathcal{T}$ with $f \in \mathcal{T} \setminus \mathcal{T}_{tr}$ in the subscript. This is because we later prove that all of our analysis can hold when training on a small fraction of tasks (Condition 4.1). Therefore, an $O(\epsilon)$ generalization error on $f \in \mathcal{T}$ can indeed reflect an $O(\epsilon)$ generalization error on $f \in \mathcal{T} \setminus \mathcal{T}_{tr}$.

Algorithm 2 Training with stochastic gradient descent (SGD)

-
- 1: **Hyperparameters:** The step size η , the number of iterations T , batch size B .
 - 2: **Initialization:** Each entry of $\mathbf{W}_O^{(0)}$ and $\mathbf{a}^{(0)}$ from $\mathcal{N}(0, \xi^2)$ and $\text{Uniform}(\{+1/\sqrt{m}, -1/\sqrt{m}\})$, respectively. $\mathbf{W}_Q$, $\mathbf{W}_K$ and $\mathbf{W}_V$ are initialized such that all diagonal entries of $\mathbf{W}_V^{(0)}$, and the first d_χ diagonal entries of $\mathbf{W}_Q^{(0)}$ and $\mathbf{W}_K^{(0)}$ are set as δ with $\delta \in (0, 0.2]$.
 - 3: **Training by SGD:** For each iteration, we independently sample $\mathbf{x}_{query} \sim \mathcal{D}$, $f \in \mathcal{T}_{tr}$ to form a batch of training prompt and labels $\{\mathbf{P}^n, z^n\}_{n \in \mathcal{B}_t}$. Each IDR pattern is sampled equally likely in each batch. For each $t = 0, 1, \dots, T-1$ and $\mathbf{W}^{(t)} \in \Psi^{(t)}$
-

$$\mathbf{W}^{(t+1)} = \mathbf{W}^{(t)} - \eta \cdot \frac{1}{B} \sum_{n \in \mathcal{B}_t} \nabla_{\mathbf{W}^{(t)}} \ell(\Psi^{(t)}; \mathbf{P}^n, z^n). \quad (4.6)$$

-
- 4: **Output:** $\mathbf{W}_O^{(T)}, \mathbf{W}_V^{(T)}, \mathbf{W}_K^{(T)}, \mathbf{W}_Q^{(T)}$.
-

where $\mathbf{P}$ is defined in (4.1). Note that the in-domain performance includes the testing performance on *unseen* tasks in $\mathcal{T} \setminus \mathcal{T}_{tr}$ that do not appear in the training samples.

Out-of-Domain Generalization Suppose the testing queries $\mathbf{x}_{query}$ follow the distribution $\mathcal{D}'$ ($\mathcal{D}' \neq \mathcal{D}$), and the binary classification tasks that map the testing queries to the labels are drawn a set $\mathcal{T}'$ ($\mathcal{T}' \neq \mathcal{T}$). Then, the *out-of-domain* generalization error can be defined as

$$\mathbb{E}_{\mathbf{x}_{query} \sim \mathcal{D}', f \in \mathcal{T}'} [\ell(\Psi; \mathbf{P}, z)]. \quad (4.5)$$

Training Algorithm The model is trained using stochastic gradient descent (SGD) with step size η with batch size B , summarized in Algorithm 2. $\mathbf{W}_Q$, $\mathbf{W}_K$ and $\mathbf{W}_V$ are initialized such that all diagonal entries of $\mathbf{W}_V^{(0)}$, and the first d_χ diagonal entries of $\mathbf{W}_Q^{(0)}$ and $\mathbf{W}_K^{(0)}$ are set as δ with $\delta \in (0, 0.2]$, and all other entries are 0. Each entry of $\mathbf{W}_O^{(0)}$ is generated from $\mathcal{N}(0, \xi^2)$, $\xi = 1/\sqrt{m}$ and each entry of $\mathbf{a}$ is uniformly sampled from $\{1/m, -1/m\}$. Besides, $\mathbf{a}$ does not update during training.

Model Pruning We also consider the case that the learned model Ψ is pruned to reduce the inference computation. Let $\mathcal{S} \subset [m]$ denote the index set of neurons in the output layer. Pruning neurons in $\mathcal{S}$ correspond to removing the corresponding rows in $\mathbf{W}_O$, resulting in the reduced matrix size of $(m - |\mathcal{S}|) \cdot m_b$.

4.3 Theoretical Characterization of In-Context Learning

4.3.1 Main Theoretical Insights

We start this section by summarizing the main theoretical insights from [464]. We consider a class of binary classification tasks where the binary labels in each task are determined by two out of M_1 *in-domain-relevant patterns*. The training data include pairs of prompt embedding and labels from a small subset of these tasks. In-domain generalization evaluates the ICL capability of the learned model on tasks using all possible combinations of these M_1 patterns. Out-of-domain generalization further evaluates the binary classification tasks that are determined by pairs of *out-of-domain-relevant patterns*, which are some linear combinations of these M_1 patterns.

(P1). Quantitative Learning Analysis With Guaranteed In- and Out-of-Domain Generalization We quantitatively prove the learned model achieves desirable generalization in both in-domain and out-of-domain tasks. The required number of training data and iterations are polynomial in β^{-1} and α^{-1} , where β represents the norm of relevant patterns, and α denotes the fraction of context inputs with the same in-domain-relevant pattern as the query. A higher α implies that the context examples offer more information about the query, consequently reducing the sample requirements and expediting the learning process.

(P2). Mechanism of Transformers in Implementing ICL We elucidate the mechanism where the learned Transformers make predictions in- and out-of-domain in context. We quantitatively show that the self-attention layer attends to context examples with relevant patterns of the query task and promotes learning of these relevant patterns. Then, the two-layer perceptron promotes the label embeddings that correspond to these examples so as to predict the label of the query accurately.

(P3). Magnitude-Based Pruning Preserves ICL We quantify the ICL generalization if neurons with the smallest magnitude after training in the MLP layer are removed and prove that the generalization is almost unaffected even when a constant fraction of neurons are removed. In contrast, the generalization error is proved to be at least $\Omega(R)$ when R fraction of neurons with large magnitude are removed.

4.3.2 The Modeling of Training Data and Tasks

In-Domain Data and Tasks Consider M_1 *in-domain-relevant (IDR)* patterns $\{\mu_j\}_{j=1}^{M_1}$ and $M_2 (= O(M_1))$ *in-domain-irrelevant (IDI)* patterns $\{\nu_k\}_{k=1}^{M_2}$ ($M_1 + M_2 = d_X$) in $\mathbb{R}^{d_X}$, where these $M_1 + M_2$ patterns are pairwise orthogonal, and

$\|\boldsymbol{\mu}_j\| = \|\mathbf{v}_k\| = \beta \geq 1$ (β is a constant) for $j \in [M_1], k \in [M_2]$. Each in-domain data $\mathbf{x}$ drawn from $\mathcal{D}$ is generated by

$$\mathbf{x} = \boldsymbol{\mu}_j + \kappa \mathbf{v}_k, \quad (4.7)$$

where $j \in [M_1]$ and $k \in [M_2]$ are arbitrarily selected. κ follows a uniform distribution $U(-K, K)$, $K \leq 1/2$. Denote $\text{IDR}(\mathbf{x}) := \boldsymbol{\mu}_j$ as the IDR pattern in data $\mathbf{x}$. Our data assumption originates from recent feature learning works on deep learning [18, 461, 630] for language and vision data. To the best of our knowledge, only [341] theoretically analyzes the performance of ICL with softmax attention, assuming all $\mathbf{x}$ are orthogonal to each other. Our assumption in (4.7) is more general than that in [341].

Each in-domain task is defined as a binary classification function that decides the label based on two IDR patterns in the query. Specifically,

Definition 4.1 (Definition of In-Domain Tasks) The in-domain task set $\mathcal{T}$ includes $M_1(M_1 - 1)$ tasks such that each task $f \in \mathcal{T}$ is defined as

$$f(\mathbf{x}) = \begin{cases} +1, & \text{IDR}(\mathbf{x}) = \boldsymbol{\mu}_a, \\ -1, & \text{IDR}(\mathbf{x}) = \boldsymbol{\mu}_b, \\ \text{random from } \{+1, -1\}, & \text{otherwise,} \end{cases} \quad (4.8)$$

where $\boldsymbol{\mu}_a, \boldsymbol{\mu}_b$ are two different patterns in $\{\boldsymbol{\mu}_j\}_{j=1}^{M_1}$ and are the decisive patterns for task f .

From (4.8), the task f outputs label +1 (or -1) if the IDR pattern is $\boldsymbol{\mu}_a$ (or $\boldsymbol{\mu}_b$). If the data contains neither of these two patterns, the label is random.

Out-of-Domain Data and Tasks Assume there are M'_1 *out-of-domain-relevant* (ODR) patterns $\{\boldsymbol{\mu}'_j\}_{j=1}^{M'_1}$ and M'_2 *out-of-domain-irrelevant* (ODI) patterns $\{\mathbf{v}'_k\}_{k=1}^{M'_2}$. Any data $\mathbf{x}$ drawn from $\mathcal{D}'$ can be generated by

$$\mathbf{x} = \boldsymbol{\mu}'_j + \kappa' \mathbf{v}'_k \quad (4.9)$$

where $j \in [M'_1]$ and $k \in [M'_2]$ are arbitrarily selected, and $\kappa' \sim U(K', K')$ for $K' = \mathcal{O}(1)$. We use $\text{ODR}(\mathbf{x}) := \boldsymbol{\mu}'_j$ to denote the ODR pattern of $\mathbf{x}$.

The set of out-of-domain tasks $\mathcal{T}'$ contains $M'_1(M'_1 - 1)$ binary classification problems that are defined in the same fashion as Definition 4.1, with the only difference of using $\{\boldsymbol{\mu}'_j\}_{j=1}^{M'_1}$ rather than $\{\boldsymbol{\mu}_j\}_{j=1}^{M_1}$ to determine labels.

Prompt Construction for Training and Testing Let l_{tr} and l_{ts} denote the length of training and testing contexts, respectively.

Training Prompt Embedding Given an input-label pair $\mathbf{x}_{query}$ and $f(\mathbf{x}_{query})$ for training, the context inputs $\mathbf{x}_i$ in $\mathbf{P}$ in (4.1) are constructed as follows. The IDR pattern is selected from $\{\mu_j\}_{j=1}^{M_1}$ following a categorical distribution parameterized by α , where $\alpha = \Theta(1) \in (0, 1]$. Specifically, each of μ_a and μ_b (the decisive patterns of task f) is selected with probability $\alpha/2$, and each of these other $M_1 - 2$ patterns elected with probability $(1-\alpha)/(M_1-2)$. The context labels are determined by task f .

Testing Prompt Embedding The context inputs for the testing query can be selected following a wide range of prompt selection methods [518, 717, 915]. Given an in-domain (or out-of-domain) task f that has decisive patterns μ_a and μ_b (or μ'_a and μ'_b), we only assume at least $\alpha'/2$ ($\alpha' \in (0, 1]$) fraction of context inputs contain the same IDR (or ODR) pattern as the query.

For the label embedding $\mathbf{y}_i$ for both training and testing, $\text{Emb}d(+1) = \mathbf{q}$, $\text{Emb}d(-1) = -\mathbf{q}$, where $\mathbf{q} \in \mathbb{R}^{d_y}$. Hence, $\mathbf{y}_i \in \{\mathbf{q}, -\mathbf{q}\}$ for $i \in [l_{tr}]$ or $i \in [l_{ts}]$.

4.3.3 In-Domain and Out-of-Domain Generalization with Sample Complexity Analysis

In order for the learned model $F(\Psi)$ to generalize all tasks in $\mathcal{T}$ through ICL, the training tasks in $\mathcal{T}_{tr}$ should uniformly cover all the possibilities of IDR patterns and labels, as stated by the following condition,

Condition 4.1 For any given $j \in [M_1]$ and either label $+1$ or -1 , the number of tasks in $\mathcal{T}_{tr}$ that map μ_j to that label is $|\mathcal{T}_{tr}|/M_1 (\geq 1)$.

Condition 4.1 is easy to meet, and $|\mathcal{T}_{tr}|$ does not have to be large. In fact, $|\mathcal{T}_{tr}|$ can be as small as M_1 . For example, let the i -th task function ($i \in [M_1 - 1]$) in $\mathcal{T}_{tr}$ map the queries with μ_i and μ_{i+1} as IDR patterns to $+1$ and -1 , respectively. The M_1 -th task function maps μ_{M_1} and μ_1 to $+1$ and -1 , respectively. We can easily verify $\mathcal{T}_{tr}$ satisfies Condition 4.1 in this case.

Following [386, 461, 755], we assume the training labels are balanced, i.e., $|\{n : z^n = +1\}| - |\{n : z^n = -1\}| = O(\sqrt{N})$. The next theorem states the training and in-domain generalization.

Theorem 4.2 (In-Domain Generalization [464]) Suppose Condition 4.1 holds. For any $\epsilon > 0$, when (i) the number of neurons in $\mathbf{W}_O$ satisfies $m \geq \Omega(M_1^2 \log M_1)$, (ii) batch size $B > \Omega(\max\{\epsilon^{-2}, M_1\} \cdot \log M_1)$, (iii) the lengths of training and testing contexts are

$$l_{tr} \geq \max\{\Omega(\log M_1/\alpha), \Omega(1/(\beta^2\alpha))\}, \quad l_{ts} \geq \alpha'^{-1}, \quad (4.10)$$

(iv) and the number of iterations satisfies

$$T = \Theta(\eta^{-1} M_1 \alpha^{-\frac{2}{3}} \beta^{-2/3} \sqrt{\log M_1}), \quad (4.11)$$

with step size $\eta \leq 1$ and $N = BT$ samples, then with a high probability, the returned model satisfies that

$$\mathbb{E}_{\mathbf{x}_{\text{query}} \sim \mathcal{D}, f \in \mathcal{T}} [\ell(\Psi; \mathbf{P}, z)] \leq O(\epsilon). \quad (4.12)$$

Theorem 4.2 characterizes the sufficient condition on the model size, the required number of iterations, and the number of prompt embedding and label pairs, such that the trained model achieves an in-domain generalization error of $O(\epsilon)$. Theorem 4.2 includes three major insights:

1. *In-domain generalization capability using a diminishing fraction of training tasks:* Because $\mathcal{T}_{\text{tr}}$ can satisfy Condition 4.1 even when $|\mathcal{T}_{\text{tr}}| = M_1$, then the number of training tasks is only a fraction $(M_1 - 1)^{-1/2}$ of the total number of in-domain tasks in $\mathcal{T}$.
2. *Context length:* The required length of training and testing contexts increase in the order of α^{-1} and α'^{-1} , respectively, which implies that a longer context is needed when the fraction of IDR patterns in the context is small.
3. *Convergence and sample complexity:* The required number of iterations and the training samples is proportional to $\alpha^{-2/3}$. This indicates that a larger fraction of the IDR pattern in the context leads to more efficient convergence and generalization.

Based on the in-domain result, we can also investigate the properties of out-of-domain generalization.

Theorem 4.3 (Out-of-Domain Generalization [464]) Suppose Condition 4.1 and conditions (i)-(iv) in Theorem 4.2 hold. For any $\mu'_1, \dots, \mu'_{M_1}, \mathbf{v}'_1, \mathbf{v}'_{M_2}$ that are pairwise orthogonal and $\|\mu'_j\| = \|\mathbf{v}'_k\| = \beta$, if

$$\mu'_j \in \left\{ \sum_{i=1}^{M_1} k_{j,i} \mu_i \mid S_j := \sum_{i=1}^{M_1} k_{j,i} \geq 1, k_{j,i} \in \mathbb{R} \right\}, \quad (4.13)$$

and $\mathbf{v}'_k \in \text{span}\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{M'_2}\}$, $j \in [M'_1]$, $k \in [M'_2]$, then with high probability, the learned model can achieve an out-of-domain generalization error of

$$\mathbb{E}_{\mathbf{x}_{\text{query}} \sim \mathcal{D}', f \in \mathcal{T}'} [\ell(\Psi; \mathbf{P}, z)] \leq O(\epsilon). \quad (4.14)$$

Theorem 4.3 indicates that a one-layer Transformer can generalize well in context, even in the presence of distribution shifts between the training and testing data. The conditions for a favorable generalization encompass the following:

- (1) the ODR patterns are linear combinations of IDR patterns with a summation of coefficients ≥ 1 , and each ODI pattern is in the subspace spanned by IDI patterns;
- (2) the testing prompt is long enough, which is linear in α'^{-1} , to include context inputs involving ODR patterns.

Remark 4.1 (Comparison with Existing ICL Analysis [341]) analyzes the generalization performance of ICL on unseen tasks under a similar data model that includes decisive and indecisive patterns. However, [341] only analyzes in-domain unseen tasks, while our results also apply to one type of out-of-domain tasks through data shift. To the best of our knowledge, only [988] studies out-of-domain generalization under the setup of linear regression problems with Gaussian inputs. They conclude that, under this setup, the covariate shift, i.e., the difference between the training and testing data distributions $\mathcal{D}$ and $\mathcal{D}'$, does not guarantee generalization. We consider classification problems under a data model different from [988]. We provide the out-of-domain generalization guarantee for one type of distribution between $\mathcal{D}$ and $\mathcal{D}'$.

4.3.4 ICL With Magnitude-Based Model Pruning

The following theorem characterizes the relation between ICL and magnitude-based model pruning.

Theorem 4.4 ([464]) *Let $\mathbf{r}_i$ be the i -row of $\mathbf{W}_O \mathbf{W}_V$, $i \in [m]$. Suppose Condition 4.1 and conditions (i)–(iv) in Theorem 4.2 hold, then there exists $\mathcal{L} \subset [m]$ with $|\mathcal{L}| = \Omega(m)$ s.t.,*

$$\begin{aligned} \|\mathbf{r}_i^{(T)}\| &\geq \Omega(1), i \in \mathcal{L}, \\ \|\mathbf{r}_i^{(T)}\| &\leq \epsilon 1/\sqrt{M_2}, i \in \mathcal{L}^c, \end{aligned} \tag{4.15}$$

where $\mathcal{L}^c$ is the complementary set of $\mathcal{L}$. Then, for any $\epsilon > 0$ and any in- or out-of-domain $\mathbf{x}_{\text{query}} \sim \mathcal{D}$ (or $\mathcal{D}'$) and corresponding $f \in \mathcal{T}$ (or $\mathcal{T}'$), pruning all neurons $i \in \mathcal{L}^c$ leads to a generalization error

$$\mathbb{E}_{\mathbf{x}_{\text{query}}, f} [\ell(\Psi_{\mathcal{L}^c}; \mathbf{P}, z)] \leq O(\epsilon + M_1^{-1/2}), \tag{4.16}$$

where $\Psi_{\mathcal{L}^c}$ represents the model weights after removing neurons in $\mathcal{L}^c$ in $\mathbf{W}_O$. In contrast, pruning $\mathcal{S} \subset \mathcal{L}$ with size $|\mathcal{S}| = Rm$, where $R \in (0, 1)$ and is a constant, and $\alpha' \geq \Omega(M_1^{-0.5})$ results in a generalization error of

$$\mathbb{E}_{\mathbf{x}_{\text{query}}, f} [\ell(\Psi_{\mathcal{S}}; \mathbf{P}, z)] \geq \Omega(R + (\alpha' M_1)^{-1}). \tag{4.17}$$

Theorem 4.4 proves that a constant fraction of neurons in $\mathcal{L}$ in the trained MLP layer has large weights, while the remaining ones in $\mathcal{L}^c$ have small weights. Pruning neurons with a smaller magnitude leads to almost the same generalization result as that of the unpruned Ψ . However, pruning neurons with a larger magnitude cause an increasing generalization error as the pruning ratio R increases. Theorem 4.4 indicates that in our setup, magnitude-based pruning on $\mathbf{W}_O$ does not hurt the model’s ICL capability.

4.3.5 The Mechanism of ICL by the Trained Transformer

Here, we provide a detailed discussion about how the generalization performance in Theorems 4.2 and 4.3 are achieved. We first introduce novel properties of the self-attention layer and the MLP layer of the learned Transformer to implement ICL.

Self-Attention Selects Contexts with the Same IDR/ODR Pattern as the Query

We first show the learned self-attention layer promotes context examples that share the same IDR/ODR pattern as the query. Specifically, for any vector $\mathbf{p} \in \mathbb{R}^{d_X+d_Y}$ that includes input $\mathbf{x}$ and the corresponding output embedding $\mathbf{y}$. We use $\text{XDR}(\mathbf{p})$ to represent the relevant pattern, which is the $\text{IDR}(\mathbf{x})$ for in-domain data and $\text{ODR}(\mathbf{x})$ for out-of-domain data. Then

Proposition 4.1 ([464]) *The trained model after being updated by T (characterized in (4.11)) iterations satisfies that, for any $(\mathbf{p}, \mathbf{W}) \in \{(\mathbf{p}_{\text{query}}, \mathbf{W}_Q^{(T)}), \{(\mathbf{p}_i, \mathbf{W}_K^{(T)})\}_{i=1}^l\}$,*

$$\|[\text{XDR}(\mathbf{p})^\top, \mathbf{0}^\top] \mathbf{W} \mathbf{p}\| \geq \Omega(\sqrt{\log M_1}), \quad (4.18)$$

$$\|[\mathbf{a}^\top, \mathbf{0}^\top] \mathbf{W} \mathbf{p}\| \leq O(\sqrt{\log M_1}(1/M_1 + 1/M_2)), \quad (4.19)$$

$$\|[\mathbf{b}^\top, \mathbf{0}^\top] \mathbf{W} \mathbf{p}\| \leq O(\sqrt{\log M_1}(1/M_1 + 1/M_2)), \quad (4.20)$$

where $\mathbf{a}$ is any IDR (or ODR) pattern that is different from $\text{XDR}(\mathbf{p})$ for in-domain (or out-of-domain) tasks, $\mathbf{b}$ is any IDI (or ODI) pattern, and $\mathbf{0}$ is an all-zero vector in $\mathbb{R}^{m_a-d_X}$.

Proposition 4.1 indicates that the self-attention layer parameters $\mathbf{W}_Q^{(T)}$ and $\mathbf{W}_K^{(T)}$ in the returned model projects $\mathbf{p}_{\text{query}}$ or context embeddings $\mathbf{p}_i$ mainly to the directions of the corresponding IDR pattern for in-domain data or ODR pattern for out-of-domain data. This can be deduced by combining (4.18), (4.19), and (4.20), since components of $\mathbf{W} \mathbf{p}$ in other directions rather than $[\text{XDR}(\mathbf{p})^\top, \mathbf{0}^\top]$ are relatively smaller. Hence, Proposition 4.1 implies that the learned $\mathbf{W}_Q^{(T)}$ and $\mathbf{W}_K^{(T)}$ remove the effect of IDI/ODI patterns. Meanwhile, (4.18) states that the $\mathbf{W}_Q^{(T)}$ and

$\mathbf{W}_K^{(T)}$ enlarge the magnitude of the IDR or ODR patterns from $\Theta(1)$ to $\Theta(\sqrt{\log M_1})$, given that the $\mathbf{W}_Q^{(0)}$ and $\mathbf{W}_K^{(0)}$ are initialized with a scalar $\delta = \Theta(1)$.

Proposition 4.1 enables us to compute the attention map of the trained model. Therefore, we have the following.

Corollary 4.1 ([464]) *For any testing query embedding $\mathbf{p}_{\text{query}} = [\mathbf{x}_{\text{query}}^\top, \mathbf{0}^\top]^\top$, let $\mathcal{N}_* \in [l]$ be the set of indices of context inputs that share the same IDR (or ODR) pattern as the in-domain (or out-of-domain) $\mathbf{x}_{\text{query}}$. Then, for any constant $C > 1$, by definition in (4.2), it holds that*

$$\sum_{s \in \mathcal{N}_*} \text{attn}(\Psi; \mathbf{P}, i) \geq 1 - \Theta(1/M_1^C). \quad (4.21)$$

Corollary 4.1 shows that after training, the attention weights become concentrated on contexts in $\mathcal{N}_*$. This means that the learned self-attention layer only selects some crucial contexts that share the same IDR/ODR pattern as the query rather than all samples uniformly or randomly.

MLP Neurons Distinguish Label Embeddings Rather Than Feature Embeddings We next show that the trained MLP layer can distinguish the label embeddings for data from different classes.

Proposition 4.2 ([464]) *Let $\mathbf{r}_i$ introduced in Theorem 4.4 be $(\mathbf{r}_{i_{d_X}}, \mathbf{r}_{i_{d_Y}})$ where $\mathbf{r}_{i_{d_X}} \in \mathbb{R}^{1 \times d_X}$, $\mathbf{r}_{i_{d_Y}} \in \mathbb{R}^{1 \times d_Y}$. Then, for any $i \in \mathcal{L}$,*

$$\mathbf{r}_{i_{d_X}}^{(T)} \bar{\boldsymbol{\mu}} / (\|\mathbf{r}_{i_{d_X}}^{(T)}\| \cdot \|\bar{\boldsymbol{\mu}}\|) \geq 1 - \Theta(1)/M_2, \quad (4.22)$$

$$\mathbf{r}_{i_{d_Y}}^{(T)} \mathbf{q}_e / (\|\mathbf{r}_{i_{d_Y}}^{(T)}\| \cdot \|\mathbf{q}_e\|) \geq 1 - \Theta(1)/M_1, \quad (4.23)$$

where $\bar{\boldsymbol{\mu}} = \sum_{k=1}^{M_1} \boldsymbol{\mu}_k^\top / M_1$, $\mathbf{q}_e = \mathbf{q}$ if $a_i > 0$ and $\mathbf{q}_e = -\mathbf{q}$ if $a_i < 0$, where a_i is the i -th entry of $\mathbf{a}$ in (4.1).

Proposition 4.2 demonstrates that neurons with indices in $\mathcal{L}$ have the following two properties. (P1) The first d_X entries of all the corresponding row vectors in $\mathbf{W}_O^{(T)} \mathbf{W}_V^{(T)}$ approximate the average of all IDR patterns $\boldsymbol{\mu}_j$, $j \in [M_1]$. (P2) The next d_Y entries of the i th row of $\mathbf{W}_O^{(T)} \mathbf{W}_V^{(T)}$ approximates the label embedding $\mathbf{q}$ when $a_i > 0$ and approximates $-\mathbf{q}$ when $a_i < 0$. (P1) indicates that the output layer focuses on all IDR patterns equally rather than any IDI pattern. (P2) indicates that the MLP layer can distinguish label embeddings for different classes.

4.4 Performance Evaluation

Experiment Setup We verify our theoretical findings using data generated as described in Sect. 4.2.1. Let $d_X = d_Y = 30$, $\beta = 3$, $K' = 5$, $K = 0.5$. The in-context binary classification error is evaluated by $\mathbb{E}_{(x,y)}[\Pr(y \cdot F(\Psi; \mathbf{P}) < 0)]$ for x following either $\mathcal{D}$ or $\mathcal{D}'$ and $\mathbf{P}$ constructed in (4.1). If not otherwise specified, we set $M_1 = 6$, $M_2 = 24$. For out-of-domain generalization, $M'_1 = 3$, $\mathbf{v}'_i = \mathbf{v}_i$ for $i \in [M'_2]$. $\boldsymbol{\mu}'_1 = 0.3 \cdot (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) + a\boldsymbol{\mu}_5 + b\boldsymbol{\mu}_6$. $\boldsymbol{\mu}'_2 = \sqrt{2}/2 \cdot (\boldsymbol{\mu}_1 + \boldsymbol{\mu}_2)$. $\boldsymbol{\mu}'_3 = \sqrt{2}/2 \cdot (\boldsymbol{\mu}_3 + \boldsymbol{\mu}_4)$. For testing, we select contexts with the two decisive patterns with $\alpha'/2$ probability each and others with $(1 - \alpha')/(M'_1 - 2)$ probability each to keep the context outputs balanced.

Model and Training The models we use include both the one-layer Transformer defined in (4.2) and the 3-layer 2-head real-world model GPT-2 [681] following [40, 909]. If not otherwise specified, we set $\alpha = 0.8$, $l_{tr} = 20$ for training. The training tasks are formulated as follows to satisfy Condition 4.1. Define $\mathbf{a}_i = \mathbf{a}_{i+M_1} = \boldsymbol{\mu}_i$ for $i \in [M_1]$, and then the $((k-1) \cdot M_1 + j)$ -th task function maps the queries with $\mathbf{a}_j$ and $\mathbf{a}_{j+k}$ as IDR patterns to $+1$ and -1 , respectively, for $j \in [M_1]$ and $k \in [U]$. For the one-layer Transformer, we use $U = 1$ and $m_a = m_b = 60$. Hence, $|\mathcal{T}_{tr}| = 6$, and there are $|\mathcal{T} \setminus \mathcal{T}_{tr}| = 24$ in-domain unseen tasks. For GPT-2, $U = 4$. Then, $|\mathcal{T}_{tr}| = 24$, $|\mathcal{T} \setminus \mathcal{T}_{tr}| = 6$. Note that we evaluate in-domain generalization error only on unseen tasks $\mathcal{T} \setminus \mathcal{T}_{tr}$, which is generally an upper bound of that defined in (4.4) after sufficient training.

Numerical Results We first verify the sufficient condition (4.13) for out-of-domain generalization. From the selection of $\boldsymbol{\mu}'$'s, we know that $S_1 = a+b$, $S_2 = S_3 = \sqrt{2}$. We vary a and b while satisfying $a^2 + b^2 + 2 \cdot 0.3^2 = 1$. Figure 4.1a shows that the out-of-domain classification error archives < 0.01 when $S_1 \geq 1$ and deviates from 0 when $S_1 < 1$, which justifies the necessity of condition (4.13). We then investigate how the context length is affected by α' , i.e., the fraction of contexts with the same IDR/ODR pattern as the query. Figure 4.1b indicates that a longer testing context length is needed when α' is smaller for in- or out-of-domain, which is consistent with the lower bound of l_{ts} in (4.10) and Theorem 4.3.

We then compare ICL with other machine learning algorithms for classification, where contexts are used as training samples for these methods. Figure 4.2a, b show that when $\alpha' = 0.8$, the advance of ICL over other algorithms is not significant, while when $\alpha' = 0.6$, ICL is the most sample-efficient for a small generalization error. Thus, ICL can remove irrelevant data and is more robust to random noise in labels than other learning algorithms.

We also investigate the effect of pruning techniques on ICL. Let $\alpha = 0.6$. Figure 4.3a shows that magnitude-based pruning does not hurt out-of-domain generalization if the pruning rate is lower than around 15%, which is the ratio of W_O neurons with a small magnitude. The generalization error increases as the pruning rate increases when pruning neurons with large weights. This is consistent with




Fig. 4.1 Out-of-domain ICL classification error on GPT-2 with (a) different S_1 on GPT-2 (b) different α' for in-domain (id) and out-of-domain (ood) generalization




Fig. 4.2 Binary classification performance of using ICL, logistic regression (Logistic), SVM with Gaussian kernel (SVM Gau.), SVM with linear kernel (SVM Lin.), 1-nearest neighbor (1-NN), and 3-nearest neighbor (3-NN) with one-layer Transformer when (a) $\alpha' = 0.8$ (b) $\alpha' = 0.6$

Theorem 4.4. Figure 4.3b justifies the impact of α' in Theorem 4.4 that larger α' can improve the performance of the pruned model.

Finally, we examine our findings regarding the mechanism of ICL in Sect. 4.3.5 using a one-layer Transformer formulated in (4.2). In Fig. 4.4a, b, we consider out-of-domain data with $a = b = 0.64$. Figure 4.4a shows that for any query $\mathbf{p}_{query}$ (or context example $\mathbf{p}_i$ for $i \in l_{ts}$), the norm of $[XDR(\mathbf{p})^\top, \mathbf{0}^\top] \mathbf{W}_Q \mathbf{p}_{query}$ (or $[XDR(\mathbf{p})^\top, \mathbf{0}^\top] \mathbf{W}_K \mathbf{p}_i$) is close to the norm of $\mathbf{W}_Q \mathbf{p}_{query}$ (or $\mathbf{W}_K \mathbf{p}_i$). This implies that the components of $\mathbf{W}_Q \mathbf{p}_{query}$ (or $\mathbf{W}_K \mathbf{p}_i$) in directions other than




Fig. 4.3 (a) Out-of-domain classification error (left y-axis for curves) with model pruning of the trained W_O using baseline (no pruning), random pruning, and magnitude-based pruning (Mag.-based), and the magnitude of each neuron of W_O (right y-axis for light blue bars) (b) Out-of-domain classification error when varying α' . These two are implemented on a one-layer Transformer




Fig. 4.4 The properties of the trained model. (a) The average norm of $W_O p_{query}$, $W_K p_i$, $[XDR(p_{query})^\top / \beta, \mathbf{0}^\top] \cdot W_O p_{query}$, and $[XDR(p_i)^\top / \beta, \mathbf{0}^\top] W_K p_i$. (b) The attention weight summation on contexts with the same ODR pattern as the query and other contexts. (c) The magnitude of the first d_X dimensions of 5 neurons in $W_O W_V$ and their angles to $\bar{\mu}$ in 400 epochs. (d) The magnitude of the rest d_Y dimensions of 10 neurons in $W_O W_V$ and their angles to q in 400 epochs. We choose 5 neurons for $a_i > 0$ and 5 for $a_i < 0$

$[XDR(p)^\top, \mathbf{0}^\top]$ are small, which is consistent with (4.19) and (4.20) in Proposition 4.1. Moreover, these norms increase from initialization during training, which justifies (4.18). Figure 4.4b depicts the concentration of attention on contexts in $\mathcal{N}_*$ after training. This verifies Corollary 4.1. Figure 4.4c, d jointly verify Proposition 4.2. The color bars represent the epochs of training. We can observe that except for some neurons, r_{id_X} grows to be close to the direction of $\bar{\mu}$ with a larger magnitude in Fig. 4.4c. Moreover, Fig. 4.4d shows for $a_i > 0$ (or $a_i < 0$), r_{id_Y} becomes close to q (or $-q$) with a large magnitude.

Part II

Advanced Topics in Foundation Models

Chapter 5

Automated Visual Prompting



Abstract Visual prompting (VP) is a parameter-efficient fine-tuning approach to adapting pre-trained vision models to solve various downstream image-classification tasks. This chapter presents **AutoVP**, an end-to-end expandable framework for automating VP design choices, along with 12 downstream image-classification tasks that can serve as a holistic VP-performance benchmark. The design space covers (1) the joint optimization of the prompts; (2) the selection of pre-trained models, including image classifiers and text-image encoders; and (3) model output mapping strategies, including nonparametric and trainable label mapping.

5.1 Introduction

Originating in the domain of natural language processing, prompting [247, 452, 751] has gained considerable popularity as a parameter-efficient fine-tuning approach for adapting pre-trained models to downstream tasks. Prompting’s methodology has recently been extended to the field of computer vision, where it is termed visual prompting (VP) [38]. In its simplest form, VP can be perceived as an in-domain model-reprogramming technique [120]. More specifically, it adjusts the inputs and outputs of a pre-trained vision model to address downstream image-classification tasks, without having to make any changes to the weights or architecture of the source model’s pre-trained backbone. As such, it stands in contrast to conventional transfer-learning methods that involve complete fine-tuning, LP (i.e., involving modifications of the trainable linear layer in the penultimate layer’s output), or zero-shot learning [674]. For instance, as illustrated in Fig. 5.1, VP adds a universal trainable data frame to the target samples at the model-input stage, and then employs a mapping function—which can be either explicitly defined or implicitly learned—to associate the source and target labels at the output stage.

While VP exhibits tremendous potential, there are two critical challenges that limit its research and development. The first is *absence of a systematic VP framework*. That is, VP design choices, such as prompts’ sizes and shapes, source models, and label-mapping (LM) strategies, have thus far only been studied separately,




Fig. 5.1 Overview and key highlights of AutoVP. The main components of AutoVP are: **Input Scaling**, which offers three initial input scale options: $\times 0.5$, $\times 1.0$, and $\times 1.5$; **Visual Prompt**, which pads the prompts to the scaled input image; **Pre-trained Classifier**, allowing users (or AutoVP) to select from four pre-trained models: ResNet18 [290], ResNeXt101-IG [565], Swin-T [538], and CLIP [674]; and **Output Label Mapping**, offering four label mapping options: Iterative Mapping (IterMap), Frequency Mapping (FreqMap), Semantic Mapping (SemanticMap), and Fully Connected Layer Mapping (FullyMap). *Bottom panel*: Given a fixed ImageNet-pre-trained classifier (ResNet18), AutoVP outperforms the state-of-the-art (ILM-VP in [108]) on all 12 different downstream image-classification tasks

generally for the purpose of delineating their distinct roles in enhancing downstream task accuracy. Ideally, such a systematic framework would automatically search for the best configurations for performance optimization. For example, [38] have demonstrated that changing the padding size of visual prompts can yield around 15% variation in final accuracy. It has also been observed that VP is better at augmenting large text-image models, such as CLIP [674], than pure vision models like ResNet50 [290]. In a study by Chen et al. [108], *iterative label mapping* (ILM) during training achieved accuracy up to 13.7% better than fixed label mapping strategies. The second critical challenge is *lack of a unified performance benchmark*: the existing literature evaluates the performance of proposed VP methods in an *ad hoc* manner, by reporting on arbitrarily selected downstream datasets, making comparisons across different methods difficult at best.

To bridge this gap, this chapter presents AutoVP proposed in [829], a solution addressing both these challenges via (1) its automated, extendable framework for joint optimization of (a) input-image scaling (i.e., prompt size), (b) visual prompts, (c) source model selection, and (d) output label-mapping strategies; and (2) its provision of a unified benchmark consisting of 12 diverse image-classification tasks with quantifiable content-similarity relative to the dataset (e.g., ImageNet) used for source model pre-training.

As shown in Fig. 5.1, the first component (i.e., input scaling) of AutoVP determines the optimal ratio between the sizes of prompts and the original images.

Table 5.1 Comparison of AutoVP with other baselines, including Linear Probing, CLIP zero-shot inference with text prompts (i.e. CLIP-TP in [674]), CLIP-VP [38], and ILM-VP [108]. The average accuracy is evaluated over 12 downstream tasks (see Sect. 5.4). For detailed information about the setting configurations, please refer to Sect. 5.3

Method	Pre-trained classifier	Prompt size	Output transformation	Output mapping number	Average accuracy (%)
Linear probing	CLIP	–	Modified last classification layer	–	79.86
CLIP-TP	CLIP	–	Fixed text prompt	1	49.54
CLIP-VP	CLIP	30	Fixed text prompt	1	76.01
ILM-VP	ResNet18 CLIP	48 30	IterMap	1	45.19 78.45
AutoVP	ResNet18	Trainable	IterMap	1/5/10	81.02
	ResNeXt101-IG		FullyMap		
	Swin-T		FreqMap		
	CLIP		SemanticMap		

The second, visual prompts, serve as trainable parameters, and undergo iterative updates during the backpropagation phase. The pre-trained model extracts pertinent features and renders predictions within the source domain; and finally, output label mapping establishes a connection between the label spaces of the source and target domains, facilitating accurate predictions in the downstream domain. The modularity of AutoVP allows for the seamless integration and easy extension of various designs for these four components.

Table 5.1 compares AutoVP against prior VP proposals and the other two baselines proposed to date: linear probing (LP) and text-prompt (TP)-based zero-shot inference. As the table shows, AutoVP is the only such framework that considers the broad range of settings that can affect VP performance. Moreover, thanks to such settings’ collective optimization, AutoVP’s configuration amounts to a breakthrough in average accuracy across 12 distinct downstream tasks. For instance, with CLIP as the pre-trained model (see Table 5.2), AutoVP’s average accuracy is 4.6% higher than CLIP-VP’s [38] and 2.1% higher than ILM-VP’s [108]. AutoVP also surpasses LP’s accuracy by 0.7% on average, suggesting that it is a competitive alternative to LP in terms of transfer learning.

5.2 Background and Related Work

Background of Visual Prompts Traditionally, to derive task-specific machine-learning models, researchers have to train or update all model parameters. But, amid the advancement of powerful foundation models, model fine-tuning and

Table 5.2 Comparison of VP testing accuracy (%) using CLIP as a pre-trained model on 12 datasets; the optimal tuning settings of AutoVP and the final prompts sizes p are also provided. In the AutoVP setting field, the notation “*Mapping- m* ” represents mapping m source classes to each target class. Bold value represents the best average performance

Dataset	AutoVP setting	AutoVP	ILM-VP	CLIP-VP	LP
SVHN [607]	FullyMap, $p = 51$	92.9 ± 0.2	91.2	88.4	65.4
CIFAR10 [423]	IterMap-1, $p = 23$	95.2 ± 0.0	94.4	94.2	95.0
Flowers102 [612]	FullyMap, $p = 16$	90.4 ± 0.6	83.7	70.3	96.9
Food101 [71]	FreqMap-1, $p = 16$	82.3 ± 0.1	79.1	78.9	84.6
UCF101 [779]	FullyMap, $p = 16$	73.1 ± 0.6	70.6	66.1	83.3
OxfordIIITPet [641]	FreqMap-10, $p = 16$	88.2 ± 0.2	85.0	85.0	89.2
CIFAR100 [423]	FullyMap, $p = 31$	77.9 ± 0.6	73.9	75.3	80.0
EuroSAT [292]	FullyMap, $p = 16$	96.8 ± 0.2	96.9	96.4	95.3
DTD [156]	FullyMap, $p = 17$	62.5 ± 0.3	63.9	57.1	74.6
ISIC [159, 830]	IterMap-1, $p = 16$	74.0 ± 1.0	73.3	75.1	71.9
FMoW [151]	FullyMap, $p = 16$	40.8 ± 0.8	36.8	32.9	36.3
GTSRB [313]	FullyMap, $p = 80$	93.1 ± 0.2	92.6	92.4	85.8
Average Accuracy		80.6	78.5	76.0	79.9

training from scratch have both become time-consuming and parameter-inefficient approaches, usually requiring large amounts of training data and storage space. To this end, VP, also known as in-domain model reprogramming, has emerged as an effective means of obtaining machine-learning models for various domain-specific tasks [120]. A well-developed pre-trained model from a source domain can be directly used to perform tasks in the target domain with little transformation of the target data. For example, we can use an ImageNet pre-trained model to classify medical images without modifying any of its parameters [827]. On the other hand, VP, along with temperature scaling, can also be used as a post-processing calibration method to align model confidence and accuracy [319, 802]. As compared to traditional approaches such as transfer learning, model fine-tuning, and training from scratch, VP is a low-complexity and model-agnostic strategy; and it is especially suitable for low-data domains.

The Design of Visual Prompts A VP framework can generally be divided into two *trainable* modules, one for *input transformation* and the other for *output transformation*. These are respectively placed at the input and output ends of a pre-trained model. In the case of input transformation, previous literature has proposed various ways to generate and place visual prompts. One of the most popular such approaches is to pad a frame around the target task image and then fill it with trainable additive-input perturbation (prompts) [38, 108, 206, 615, 827, 908]. Next, since the output logits of the source pre-trained model are still in the source domain, further output transformation (e.g., LM) is required to obtain the target-domain logits. One naive way of achieving this is randomly mapping (RandMap) m source labels onto the target labels. Tsai et al. [827] found that frequency-based LM

(FreqMap), which constructs its LM from the source-label prediction distribution of the target-domain data, can further improve the accuracy of downstream tasks. However, [108] argued that FreqMap lacks interpretability and that its interaction with VP is difficult to measure. To address that problem, the authors proposed iterative LM (IterMap), a transformation of FreqMap that enables it to concurrently design LM and visual prompts. Yang et al. [945], meanwhile, proposed a semantics-based LM approach that aligns source and target classes that have similar semantic meanings. And [496] utilized a prototypical verbalizer to map a mask token to downstream labels, thus providing a different perspective on LM. In this paper, we follow a similar design to [38], in which visual prompts are placed around images for input transformations, and there are four mapping methods for output transformations. Further details will be presented in Sect. 5.3.

Non-universal Visual Prompts Instead of utilizing universal input prompts, some researchers have focused on designing input-aware prompting models [1037, 1038]. For instance, [107] generated class-wise visual prompts to improve model robustness. Similarly, to address accuracy drops caused by low-voltage-induced bit errors, [790] proposed an input-aware add-on module to generate a robust prompt; and [544] proposed the Prompt Generation Network (PGN), which generates visual prompt token vectors based on input images, allowing for more adaptive and context-aware prompting.

Although input prompting is commonly applied directly to the target image, researchers have also developed other prompting methods, such as *convolutional visual prompt* [828], which learns prompting parameters in a small convolutional structure through self-supervision tasks without knowledge of ground truths, and *visual prompt tuning* [371, 770], which learns prompting parameters at intermediate layers of a source model. In this chapter, we focus on a pixel-level VP approach using a task-specific prompting model for each image-classification task. As such, our approach closely resembles real-world scenarios in which a pre-trained source model remains unmodified, and external variations are not introduced internally.

Visual Prompting Meets Differential Privacy In [30, 490], the authors show that VP with a pre-trained model (trained on non-private data) can improve the privacy-accuracy tradeoff in off-the-shelf training mechanisms of differential privacy (DP). Particularly, when VP is used in PATE (Private Aggregation of Teacher Ensembles) [639], a DP training mechanism, [490] shows that the classification accuracy under a privacy constraint can achieve the current state-of-the-art performance on some common benchmark of image classification tasks.

5.3 AutoVP Framework

Following the system overview of AutoVP in Fig. 5.1, we present its four main components (Input Scaling, Visual Prompt, Pre-trained Classifier, and Output

Label Mapping) and its hyper-parameter tuning feature, which enables the joint optimization of these components. Our framework can be extended to support user-defined configurations.

5.3.1 Input Scaling

In our implementation of AutoVP, we choose frame-shape prompts as the default prompting template. Hence, the visual prompt sizes p represent the width of the frame, and its actual number of parameters is $2cp(h + w - 2p)$, where c , w , and h are color channels, width, and height respectively. Although the input image size is determined by the source model, when fine-tuning to a downstream dataset from a source model, there is design freedom to resize the target images and use the remaining space for visual prompts. For instance, if the source model takes images with size 224×224 as input, one can scale the target image size to 128×128 , resulting in the final visual prompt of size $p = (224 - 128)/2 = 48$. It was shown in [38] and [908] that the prompt size (p) plays a key role in VP performance. To automate the process of optimizing image resizing scale, we design a gradient-based optimization algorithm to implement the *input scaling module*, which is implemented using `kornia.geometry.transform()` from the Kornia library [703]. The `transform()` function integrates a range of geometric transformations for 2D images into deep learning, including differentiable image scaling. In addition to image resizing, the prompt size p will also scale up or down to fill the remaining space. Furthermore, to facilitate the optimization of image resizing and avoid bad local minima, we set the default image size to 128 along with three initial scales: 0.5, 1.0, and 1.5 to optimize, and the corresponding prompt sizes p are 80, 48, and 16 respectively. Consequently, the input scaling module allows AutoVP to obtain the optimal image resizing scale and prompt size (p).

5.3.2 Visual Prompt

For the visual prompt module, AutoVP adds universal pixel-level prompts around all (resized) input images. Let $x_t \in \mathbb{R}^{N_t}$ denote the *target* (flattened) input image (of N_t -dimension), $\tilde{x}_t \in \mathbb{R}^{N_s}$ be the prompted image, which fits the input dimension (N_s) of the pre-trained source model f_{θ_s} (θ_s denotes its weights), $\delta \in \mathbb{R}^{N_s}$ be a trainable universal perturbation, and $\mathcal{M}_p \in \{0, 1\}^{N_s}$ be a binary mask of prompt size p , indicating the prompting area. Hence, the prompted image $\tilde{x}_t$ can be formulated as:

$$\tilde{x}_t = \mathcal{P}(x_t) = \text{InputScaling}_p(x_t) + \underbrace{\mathcal{M}_p \odot \sigma(\delta)}_{\text{Prompts}}. \quad (5.1)$$

The prompts are initialized as 0 and formally defined as $\mathcal{M}_p \odot \sigma(\delta)$, where σ is the Sigmoid function that maps the input to a value between 0 and 1 (the scaled input pixel value range), ensuring it has the same numerical range as the input image. We then update δ using gradient descent algorithms.

5.3.3 Pre-trained Classifier

After applying the prompts to the resized image through the preceding stages, the prompted image is subsequently fed into the pre-trained model, which serves as a feature extractor to generate predictions in the source domain. We include four representative pre-trained models in our AutoVP framework: ResNet18 [290], ResNeXt101-IG [565], Swin-T [538], and a vision-language multi-modal model, CLIP [674] with the ViT-B/32 vision encoder backbone. Note that in AutoVP, the weights of the pre-trained classifiers are frozen and kept unchanged.

5.3.4 Output Label Mapping

The pre-trained models predict target data to source labels, while the last mile for VP is to map predictions on the source labels to target classes. As illustrated in Fig. 5.1, AutoVP provides four types of output mapping, and they can be generally categorized into two groups. (i) *nonparametric* label mapping: frequency mapping (FreqMap) and semantic mapping (SemanticMap), which are defined during the initialization of VP training and remain unchanged throughout the training process; and (ii) *trainable* label mapping: iterative label mapping (IterMap) and fully connected layer mapping (FullyMap). These two methods dynamically adjust the mapping based on the prompted images.

AutoVP incorporates four output label mappings: frequency mapping (FreqMap), iterative mapping (IterMap), semantic mapping (SemanticMap), and fully connected layer mapping (FullyMap). Figure 5.2 illustrates different methods. In the following, we provide more details of each mapping algorithm.

- **Frequency Mapping (FreqMap)** is proposed in [827]. It utilizes the source-label prediction distribution of the target-domain data to map each target class to the top- m most frequent source classes. Let $\mathcal{Y}_s = \{0, \dots, K_s - 1\}$ and $\mathcal{Y}_t = \{0, \dots, K_t - 1\}$ be the set of source and target labels, where K_s/K_t are the numbers of source/target classes. Consider $\tilde{\mathcal{X}}_t$ collects all prompted images of label y_t in target domain $\mathcal{D}_t$, i.e. $\tilde{\mathcal{X}}_t = \{\tilde{x}_{t_i} = \mathcal{P}(x_{t_i}) | (y_{t_i} = y_t), (x_{t_i}, y_{t_i}) \in \mathcal{D}_t\}$, then when $m = 1$, the mapping of y_t can be defined as:

$$y_t \leftarrow y_s^* = \arg \max_{y_s \in \mathcal{Y}_s} \left(\sum_{\tilde{x}_t \in \tilde{\mathcal{X}}_t} \text{Pred}(f_{\theta_s}(\tilde{x}_t), y_s) \right), \quad (5.2)$$




Fig. 5.2 Illustration of four output label mapping methods

where

$$\text{Pred}(f_{\theta_s}(\tilde{x}_t), y_s) = \begin{cases} 1, & \text{if } y_s = \arg \max f_{\theta_s}(\tilde{x}_t) \\ 0, & \text{otherwise} \end{cases} \quad (5.3)$$

The objective of FreqMap is to map the target label y_t to the source label y_s^* , which is the most frequent label that f_{θ_s} classified $\tilde{x}_t$ as. If a source class is selected as the most frequently predicted class for multiple target classes, it will be assigned to the target class that has the highest count of predictions. The general many-to-one frequency mapping algorithm is specified in [829]. Moreover, *random label mapping* (RandMap) can be viewed as a special case of FreqMap by randomly assigning a subset of source labels to a target label.

- **Iterative Mapping (IterMap, or ILM)** is proposed in [108], which is an iterative approach for updating FreqMap. IterMap performs the frequency mapping at the beginning of each training epoch to obtain a new mapping distribution that aligns with the updated prompts.
- **Semantic Mapping (SemanticMap)** follows the works from [945] and [951]. We utilize the text encoder of CLIP to generate the embeddings of the names of the source and target classes. Subsequently, we map the source-target pairs based on the highest cosine similarity score between their respective embeddings. Hence, SemanticMap can be utilized in any of the three vision pre-trained models (ResNet18, Swin-T, and ResNeXt101-IG) by establishing mappings between the target classes and semantically similar classes from ImageNet-1K. However, SemanticMap is not applicable for CLIP, as it does not have an explicit set of source domain classes.




Fig. 5.3 Hyper-Parameter Tuning Selection. Illustration of the end-to-end hyper-parameter tuning process in AutoVP with a total of 222 possible configurations

- **Fully Connected Layer Mapping (FullyMap)** uses a linear layer to map the source output logits to target classes [31]. FullyMap can be represented as $L_t = w \cdot L_s + b$, where L_s is the output logits from the source pre-trained model, w and b are the weight and bias vector of the linear layer, and L_t is the output of the linear layer which also serves as the final output logits of the VP model.

5.3.5 End-to-End Hyper-Parameter Tuning

Figure 5.3 presents the stages involved in the tuning process, while the **Visual Prompt** component depicted in Fig. 5.1 is not involved in the tuning process, as it does not contain any hyper-parameters. During the **Input Scaling** stage, the initial scale of the input image is determined, and users can choose whether to learn the resize scale during training. In the **Pre-trained Classifier** stage, users have the option to select from four pre-trained models to serve as the feature extractor. The **Output Label Mapping** stage offers four mapping methods to choose from. For FreqMap, IterMap, and SemanticMap, users can specify the number of source classes that are mapped to a single target class.

Given its flexibility and modularity, its users must consider numerous settings ($n = 222$), including how big the initial input image should be, whether to use a trainable resizing module, which pre-trained classifiers to adopt, what output-mapping method to implement, and the number of source labels to map for each target label. To speed up the tuning operation and save computational resources, we use Ray Tune [498] along with an early-stop strategy for terminating poor trails. In our experiments, we employed grid searches to test all configurations. An ASHA scheduler [473] was used to retain the top- n tasks, and we continued training them while stopping the remaining tasks early. We established experimentally that $n = 2$ top tasks were enough to obtain the optimal setting. When the few-epoch tuning process (training 2–5 epochs with each setting) is complete, we select the setting

having the highest testing accuracy and conduct complete training using that setting. By using hyper-parameter tuning, AutoVP can efficiently find the best configuration of VP and lead to significant accuracy improvement in downstream tasks.

5.4 Performance Evaluation

We evaluated the performance of AutoVP on 12 downstream datasets (CIFAR10, CIFAR100, ISIC, SVHN, GTSRB, Flowers102, DTD, Food101, EuroSAT, OxfordIIITPet, UCF101, and FMoW), which are common datasets when measuring transfer learning generalization. Detailed descriptions of these datasets are given in [829]. We repeated each AutoVP experiment in triplicate, utilizing a learning rate of 40 with the SGD optimizer for CLIP, and a learning rate of 0.001 with the Adam optimizer for the other pre-trained models. The results of the baselines (CLIP-VP [38] and ILM-VP [108]) were extracted from the reported accuracies in their respective papers.

Comparison of AutoVP and Prior Work To ensure that our comparison of AutoVP against previously proposed VP approaches was fair, we fixed its source model but relaxed its other hyperparameter tunings. The results of using CLIP as the source model are presented in Table 5.2, along with the optimal settings arrived at. We compared AutoVP against LP and two state-of-the-art VP baselines, CLIP-VP and ILM-VP, whose configurations can also be found in Table 5.1. With the optimal configuration chosen via the tuning process, AutoVP outperformed these other approaches by up to 6.7% on nine of the 12 target datasets. Additionally, AutoVP surpassed the LP baseline on half those datasets, by a maximum of 27.5% in the case of SVHN. AutoVP also obtained the best average accuracy.

We observed that AutoVP employed FullyMap as the output transformation on most datasets. We speculate that this might have been because the linear layer has more parameters and thus allows the achievement of better results. Also, when AutoVP selected initial image scales, it had a tendency to scale up those images with relatively small prompt sizes. This allowed the VP model to allocate more attention to the image itself, leading to improved overall performance. As shown in Fig. 5.1, when ResNet18 was used as the source model, AutoVP outperformed ILM-VP by 24.8% on average.

AutoVP with Source Model Selection We also allowed AutoVP to search the optimal source model for downstream tasks. The optimal settings selected by AutoVP, and a comparison of experimental results can be found in Table 5.3. Our experimental results show that Swin-T was the pre-trained model most frequently chosen by AutoVP as most suitable, i.e., in the cases of eight of the 12 datasets. On average, this choice resulted in 0.43% better accuracy than when CLIP was utilized as the fixed pre-trained backbone. On the DTD and Flowers102 datasets, however, Swin-T’s performance was better than CLIP’s by much more: i.e., 6.80

Table 5.3 AutoVP with source model selection. This table displays the best tuning setting without any restriction on the choice of pre-trained model, and shows the test accuracy (%) of AutoVP and the LP baseline of the chosen model across 12 datasets. Bold value represents the best average performance

Dataset	AutoVP setting	AutoVP	LP
SVHN	CLIP, FullyMap, $p = 51$	92.86 ± 0.18	65.40
CIFAR10	ResNeXt101-IG, FullyMap, $p = 48$	95.89 ± 0.07	93.89
Flowers102	Swin-T, FullyMap, $p = 16$	93.48 ± 0.52	95.75
Food101	CLIP, FreqMap-1, $p = 16$	82.28 ± 0.09	84.60
UCF101	Swin-T, FullyMap, $p = 16$	72.96 ± 0.26	75.96
OxfordIIITPet	Swin-T, FullyMap, $p = 16$	90.20 ± 0.55	93.04
CIFAR100	ResNeXt101-IG, FullyMap, $p = 48$	79.76 ± 0.47	76.09
EuroSAT	Swin-T, FullyMap, $p = 16$	95.98 ± 0.02	95.50
DTD	Swin-T, FullyMap, $p = 16$	69.25 ± 0.58	71.49
ISIC	Swin-T, FullyMap, $p = 16$	71.66 ± 1.45	72.22
FMoW	Swin-T, FullyMap, $p = 48$	39.79 ± 0.83	32.73
GTSRB	Swin-T, FullyMap, $p = 55$	88.10 ± 2.11	74.97
Average accuracy		81.02	77.64

and 3.08%, respectively. These findings highlight how multiple pre-trained models can be leveraged to enhance performance across a diverse range of datasets.

Data Scalability To understand how AutoVP would perform in a data-limited regime, we gradually and uniformly reduced the amount of training data to 50%, then 25%, then 10%, and finally 1% of each training dataset’s original size. The experimental results in Fig. 5.4 indicate that AutoVP consistently outperformed LP across all 12 datasets, and that its relative performance was especially high in the two scenarios with the lowest data volumes, i.e., 10 and 1% data usage.

Performance Evaluation on ID/OOD Downstream Tasks We evaluate the out-of-distribution (OOD) extent of each dataset relative to the pre-trained CLIP by considering the average confidence score [275] and the CLIP zero-shot inference. The accuracy gains achieved through VP (Fig. 5.5) were computed as the difference in accuracy between AutoVP and LP or non-VP approaches (i.e. visual prompts were removed and output mapping retained). We observed that the datasets that were more in-distribution (ID), with higher confidence scores and higher zero-shot accuracy, exhibited smaller accuracy gains from VP. Conversely, datasets that were more OOD, characterized by lower confidence scores and lower zero-shot accuracy, had their accuracy improved more through AutoVP.




Fig. 5.4 Data Scalability. The chart presents the average accuracy of AutoVP and LP across the 12 datasets with varying data percentages: 100, 50, 25, 10, and 1%. The green bar represents the accuracy gains achieved by AutoVP compared to Linear Probing (LP)




Fig. 5.5 Accuracy Gains with CLIP. The right side of the chart indicates a higher out-of-distribution (OOD) extent, accompanied by larger gain values. Conversely, the left side shows lower gain values

Chapter 6

Prompting Large Language Models with Privacy



Abstract Numerous studies have highlighted the privacy risks associated with pretrained large language models. This chapter offers a unique perspective by demonstrating that pretrained large language models can effectively contribute to privacy preservation. We present a locally differentially private mechanism called DP-Prompt proposed in Utpala et al. (Locally differentially private document generation using zero shot prompting. In: Findings of the Association for Computational Linguistics: EMNLP 2023, pp. 8442–8457. Association for Computational Linguistics (2023)), which leverages the power of pretrained large language models and zero-shot prompting to counter author de-anonymization attacks while minimizing the impact on downstream utility.

6.1 Introduction

The vast amount of online text data has the potential to reveal numerous user attributes, making individuals easily identifiable [315, 662, 694]. While private information can be directly disclosed through specific phrases in the text, it can also be implicitly inferred. For instance, linguistic patterns embedded within the text can inadvertently facilitate authorship attribution [392, 759], leading to unintended privacy leakage. An illustrative real-world scenario is the AOL search data leak in August 2006 [642]. The incident unfolded when AOL mistakenly released detailed search logs of their users, wrongly assuming that the data had been adequately anonymized through the use of random user IDs. Unfortunately, the released logs contained sufficient personally identifiable information, leading to the identification of numerous individuals [49, 381]. This breach of privacy triggered widespread public outcry and led to the initiation of class action lawsuits.

This case is just one among many that highlights the limitations of ad-hoc privacy approaches that may give the impression of providing privacy but ultimately fall short. Differential privacy (DP) provides a rigorous treatment for the notion of data privacy by providing plausible deniability by precisely quantifying the deviation in the model's output distribution under modification of a small number of data points [202, 203]. The provable guarantees offered by DP, coupled with its compelling

properties such as immunity to arbitrary post-processing and graceful composability, have established it as the de facto standard for privacy. DP has witnessed widespread adoption and numerous deployments in both private [29, 207, 605] and public organizations [9].

To address the issue of deanonymization attacks, various approaches have been proposed within the DP framework. These approaches encompass word-level strategies [97, 224, 931] where noise is added at the word level, as well as sentence-level techniques [576] where noise is added at the sentence level. However, recent research by Mattern et al. [572] has identified limitations in word-level approaches, particularly their disregard for contextual information. To overcome these limitations, Mattern introduced a mechanism that fine-tunes the GPT-2 model [680] specifically for paraphrasing tasks, resulting in the generation of sanitized versions of documents. While promising, the approach is limited by their reliance on annotated paraphrasing data, extensive computing resources for larger models, and the quality of annotations.

This chapter presents the **DP-Prompt** approach proposed in [841], a novel and straightforward solution to address deanonymization attacks. DP-Prompt leverages pretrained large language models by directly prompting them to generate paraphrases. These paraphrases are then released as sanitized documents in a zero-shot manner (see Fig. 6.1). The motivation for this approach stems from two important factors. Firstly, recent research [58, 572] has shown that paraphrasing is a robust defense mechanism against deanonymization attacks. Secondly, growing evidence suggests that pretrained large language models can effectively tackle complex tasks without the need for task-specific and expensive fine-tuning [81, 149, 155, 417, 620], through zero-shot prompting. By harnessing the capabilities of pretrained large language models (LLMs), DP-Prompt offers a straightforward and powerful solution to mitigate the risk of deanonymization. It provides a promising alternative that can be widely applicable, particularly in the context of on-device large language models where text completion tasks require significantly fewer resources.

6.2 Background and Related Work

6.2.1 Differential Privacy (DP)

A mechanism $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{V}$ achieves ϵ -PureDP if, for all inputs $D, D' \in \mathcal{D}$ that differ in one element, and for all $V \subseteq \text{Range}(\mathcal{M})$, $\Pr[\mathcal{M}(D) \in V] \leq \exp(\epsilon) \Pr[\mathcal{M}(D') \in V]$ [202].

Metric Differential Privacy (Metric-DP) [25, 106] is a relaxation of Pure-DP that applies to data represented in a general metric space. For a given distance metric $d : \mathcal{D} \times \mathcal{D} \rightarrow \mathbb{R}_+$, a mechanism $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{V}$ achieves ϵd -MetricDP if, for any $D, D' \in \mathcal{D}$ and for all $V \subseteq \text{Range}(\mathcal{M})$, $\Pr[\mathcal{M}(D) \in V] \leq \exp(d(D, D')) \Pr[\mathcal{M}(D') \in V]$.




Fig. 6.1 Overview of the DP-Prompt mechanism [841]. Given a private document (D), DP-Prompt generates a sanitized document (P) while ensuring local differential privacy. The level of privacy guarantee is controlled by adjusting the sampling temperature (T) during the decoding process

Local differential privacy (LDP) [199, 388, 922] is a privacy framework where data is locally perturbed before transmission, considering the presence of an untrusted data collector or server. The formal definition of LDP is as follows:

Definition 6.1 (PureLDP) A randomized mechanism $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{V}$ is said to be ϵ -PureLDP if for any pair of inputs $D, D' \in \mathcal{D}$ and for all $V \subseteq \text{Range}(\mathcal{M})$

$$\Pr[\mathcal{M}(D) \in V] \leq \exp(\epsilon) \Pr[\mathcal{M}(D') \in V].$$

There is a growing consensus that, despite the assurance of formal guarantees, it is imperative to subject differentially private mechanisms to robust privacy attacks that simulate strong and malicious adversaries [65, 363]. Such evaluation allows to effectively assess the empirical privacy provided by the mechanism in real-world scenario. To this end we define four attack models depending its adaptivity and mode of access.

Definition 6.2 (Attack Models) Consider a collection of private documents $(D_1, \dots, D_n)$ from distribution $\mathcal{D}$ with associated author identities $(a_1, \dots, a_n)$ and embeddings $(E_1, \dots, E_n) \sim \mathcal{E}$.

For text-to-text sanitization using mechanism $\mathcal{M}_{\text{text}}$, the sanitized documents are represented as $(P_1, \dots, P_n) \sim \mathcal{P}_{\mathcal{M}_{\text{text}}}$. For text-to-embedding sanitization via mechanism $\mathcal{M}_{\text{embedding}}$, the sanitized embeddings are denoted as $(N_1, \dots, N_n) \sim \mathcal{N}_{\mathcal{M}_{\text{embedding}}}$.

- **Static Attacker with Embedding Access:** Has access to clean documents $(D_1, \dots, D_n)$ but lacks access to sanitized versions $(P_1, \dots, P_n)$.
- **Static Attacker with Text Access:** Doesn't have access to sanitized embeddings $(N_1, \dots, N_n)$ but only to the clean embeddings $(E_1, \dots, E_n)$.
- **Adaptive Attacker with Embedding Access:** Has access to sanitized embeddings $(N_1, \dots, N_n)$. Hence, trains a de-anonymization model to adapt to the DP mechanism $\mathcal{M}_{\text{embedding}}$.
- **Adaptive Attacker with Text Access:** Has access to sanitized text $(P_1, \dots, P_n)$. Consequently, trains a de-anonymization model to adapt to the DP mechanism $\mathcal{M}_{\text{text}}$.

It is important to note that the adaptive attacker is a more formidable adversary since it adapts to the characteristics of the mechanism $\mathcal{M}$, whereas the static attacker only has access to clean documents/clean embeddings without any added noise. The mode of access—either raw text or abstracted embeddings—offers further nuances, determining the exact nature of the data an attacker can exploit.

6.2.2 Document Processing with Privacy

The previous work on releasing private documents can be categorized into three approaches based on the level at which noise is added (see Table 6.1 for concise summary). These approaches are:

Table 6.1 Comparison of different privacy methods. The “Privacy level” indicates the privacy guarantee provided by each mechanism. “Fine-tuning” denotes whether the mechanism involves fine-tuning a model as an intermediate step. The last column, “Generates sanitized doc,” indicates whether the mechanism can output a fully sanitized document instead of just sanitized embeddings

Mechanism	Privacy level	Requires fine-tuning	Generates sanitized doc
Madlib [224]	Word level metric-DP	No	Yes
Mahanolbis [931]	Word level metric-DP	No	Yes
TEM [97]	Word level metric-DP	No	Yes
Truncated laplace [576]	Sentence level pure-DP	No	No
Deep candidate [576]	Sentence level pure-DP	Yes	No
Paraphraser [572]	Document level pure-LDP	Yes	Yes
DP-prompt [841]	Document level pure-LDP	No	Yes

Word-Level Approaches MadLib [224] is a word-level mechanism that applies Laplace noise to word embeddings and maps them to the nearest word in the vocabulary, demonstrating the differential privacy (DP) guarantee of MadLib under the Euclidean metric. An extension of this approach involves using a regularized Mahalanobis metric instead [931]. In contrast, the TEM mechanism utilizes the exponential mechanism to transform the privatization step into a selection problem [97]. Furthermore, there is a recent development known as CusText [126], which focuses on developing customized mapping mechanisms for each individual word in the vocabulary [126]. All of these approaches are word-level mechanisms and have been shown to have significant limitations, such as their disregard for contextual information [572].

Sentence-Level Approaches Sentence-level mechanisms based on Sentence Transformer [699] were introduced in [576]. They proposed two approaches: one approach where noise is added to sentence embeddings, and another more complicated approach based on maximizing Tukey depth [255, 834].

Document-Level Approaches A document-level Local Differential Privacy (LDP) mechanism was introduced, where GPT-2 is fine-tuned for a paraphrasing task [572]. Our approach, DP-Prompt, draws inspiration from their work, but instead of resource-intensive fine-tuning, we use a zero-shot approach with pretrained models for efficient and effective generation of sanitized documents. Furthermore, the recently proposed DP-BART [346] employs BART [454], an encoder-decoder model. In DP-BART, noise is added to the encoder’s output, and the decoder is fine-tuned to adapt to this noisy encoder output.

Adversarial Methods Parallel to differentially private approaches, other techniques have been proposed that utilize Adversarial Learning [673, 750] and Data Poisoning [377, 879]. However, these methods typically require access to a surrogate classifier. In contrast, our method is zero-shot, requiring neither fine-tuning nor access to a classifier.

Differentially Private Training/Fine Tuning There is extensive research on differentially private training or fine-tuning of language models [27, 391, 480, 571, 961]. They aim to make language models resistant to various kinds of data leakage attacks [46, 90, 92, 178]. It is important to emphasize that this line of work is completely distinct from DP-Prompt [841], as it focuses on training language models on private data, while the goal of DP-Prompt is to generate sanitized documents from private documents using pretrained language models.

6.3 DP-Prompt

Language models use a decoder network to generate sequential text output. Given a context or prompt represented by a sequence of tokens $C = (c_1, \dots, c_m)$, the language model generates text by sampling tokens from a conditional distribution

Algorithm 3 DP-prompt

```

1: Input: language model (LM), private document (D), prompt template ( $T$ ), clipping vector
    $\mathbf{b} \in \mathbb{R}^{|\mathcal{V}|}$ , temperature  $T \in \mathbb{R}_+$ , paraphrase tokens  $n$ .
2: Output: Sanitized Doc (P)
3:  $P \leftarrow []$ ,  $C \leftarrow \text{GeneratePrompt}(D, T)$ 
4:  $C_{\text{tokens}} \leftarrow \text{Tokenize}(C)$ 
5: for  $i \leftarrow 1$  to  $n$  do
6:    $\mathbf{u} \leftarrow \text{LM}(C_{\text{tokens}})$ 
7:    $\mathbf{u}' \leftarrow \text{ClipAndScale}(\mathbf{u}, \mathbf{b}, T)$ 
8:    $\mathbf{p} \leftarrow \text{ConvertToProbabilities}(\mathbf{u}')$ 
9:    $v \leftarrow \text{SampleTokens}(\mathbf{p})$ 
10:   $P \leftarrow P \cup [v]$ ,  $C_{\text{tokens}} \leftarrow C_{\text{tokens}} \cup [v]$ 
11: end for
12:  $P \leftarrow \text{Detokenize}(P)$ 
13: Return: P

```

$\Pr_{|C}(x_1, \dots, x_n) = \prod_{i=1}^n \Pr_{|C}(x_i | x_1, \dots, x_{i-1})$. In this distribution, the logits $\mathbf{u} \in \mathbb{R}^{|\mathcal{V}|}$ are transformed using the softmax function with a temperature T , where $p_{ij} = \frac{\exp(\frac{u_{ij}}{T})}{\sum_{j=1}^{|\mathcal{V}|} \exp(\frac{u_{ij}}{T})}$, and $\mathcal{V}$ represents the vocabulary.

This process of sequentially generating text can be regarded as a problem of selecting tokens at each step. Hence, to make the generation step differentially private, one must replace it with a differentially private version of the selection process. One commonly used and well-known differentially private mechanism is the exponential mechanism [575], which is defined as follows:

Definition 6.3 (Exponential Mechanism) Given an utility function $u : \mathcal{D} \times \mathcal{V} \rightarrow \mathbb{R}$. The exponential mechanism $\mathcal{M}_{\text{Exp}} : \mathcal{D} \rightarrow \mathcal{V}$ is a randomized algorithm with output distribution $P[\mathcal{M}_{\text{Exp}}(D) = v] \propto \exp\left(\frac{\epsilon u(D, v)}{2\Delta u}\right)$, where $\Delta u = \max_{D, D', v} |u(D, v) - u(D', v)|$ is sensitivity.

In our case, the utility of token $v_j \in \mathcal{V}$ at each step i is simply the logit $u_{ij} \in \mathbb{R}$. Hence, one can make text generation differentially private using the exponential mechanism.

Extensive research has shown that paraphrasing documents helps conceal author identity [58, 572, 694]. Considering recent advancements where tasks are formulated as prompts and language models are tasked to complete them [81, 685, 885], we directly prompt the language model to generate paraphrases. Therefore, given a private document D and a specific prompt template instructing the language model to generate a paraphrase, such as $T := \text{"Paraphrase of the document:"}$ we combine D and T to create a context C . By utilizing this context, we execute the text generation procedure in a differentially private manner to produce a paraphrase. We refer to this procedure as DP-Prompt. Algorithm 3 outlines the specific steps of the proposed DP-Prompt method in [841]. The formal guarantee of achieving ϵ -PureLDP is provided by the following theorem:

Theorem 6.1 ([841]) *Suppose the language model has not been pretrained on the private documents distribution $\mathcal{D}$. If the final logits $\mathbf{u} \in \mathbb{R}^{|\mathcal{V}|}$ satisfy the condition $b_1 \leq u_i \leq b_2, \forall i$, and the DP-Prompt run with a temperature T for generating n tokens, then it can be proven that the generated output satisfies $(2n(b_2 - b_1)/T)$ -LDP.*

6.4 Performance Evaluation

Experiment Setup Note that we are comparing DP-mechanisms with different levels of differential privacy. Therefore, in our experiments, we focus on evaluating the empirical privacy rather than the theoretical privacy(ϵ) for effective and realistic assessment. As a result, we plot the author identification F1 score, which is calculated by conducting de-anonymization attacks on the sanitized documents. This score indicates the potential for privacy breaches. On the other hand, the y-axis represents the sentiment F1 score, which measures the utility of the sanitized documents.

Datasets We conduct experiments using IMDB movie reviews and Yelp business reviews, both of which contain author and sentiment labels. The IMDB dataset has a size of 15,000, while the Yelp dataset has 17,336 samples. For both datasets, sentiment analysis is a 2-class classification task, and the author identification task is a 10-class classification task.

Implementation Details For the embedding-level attacker, we utilize 3-Layer MLPs with ReLU activation functions and train them on sentence embeddings [699]. For the text-level attacker, we fine-tune BERT [182]. Regarding the static attacker, the clean set of documents is used for training and validation, while the sanitized documents serve as the test set. On the other hand, for the adaptive attacker, all three sets (training, validation, and testing) consist of sanitized documents.

Baselines

- For each of word level mechanisms, (Madlib [224], Mahalanobis [931], TEM [97]) we run the mechanisms for 8 ϵ 's given $\epsilon = \{2, 5, 8, 11, 14, 17, 20, 25\}$
- For each of sentence level mechanisms (Truncated-Laplace [576], Deep-Candidate [576]), we run the mechanisms for 11 ϵ values given by $\epsilon = \{5, 10, 20, 30, 40, 50, 75, 100, 150, 200\}$.
- For Paraphraser [572] and DP-Prompt with open source models we run decoding at 5 temperatures $\{0.75, 1.0, 1.25, 1.5, 1.75\}$. For DP-Prompt we run ChatGPT at temperatures $\{1.0, 1.25, 1.5, 1.75, 2.0\}$.

Further, we also consider F1 scores on Clean (without noise added) embeddings/documents and performance of uniformly a random classifier.

Numerical Results We compare 6 baselines (Madlib, Mahalanobis, Tem, Truncated-laplace, Deep-candidate, Paraphraser) run with configurations above with DP-Prompt on ChatGPT (GPT-3.5). Except for DP-Prompt, we run each mechanism to 3 times to produce 3 different sanitized documents and plot mean author F1 identification score on x-axis and show 2σ band around mean sentiment F1 score. Results are show in Fig. 6.2.

The results clearly demonstrate the superior performance of DP-Prompt with ChatGPT (GPT-3.5). Notably, DP-Prompt exhibits significantly higher utility on the y-axis for a chosen empirical privacy value on the x-axis. All word-level mechanisms show a similar privacy-utility tradeoff. Regarding sentence-level mechanisms, the truncated Laplace mechanism performs decently, while in the static attack experiments, Deep-candidate is reduced to a random classifier due to the distribution shift caused by sentence recoding.

Furthermore, in the case of clean reviews (i.e., without any noise), the embedding-level attacker can accurately identify the author among 10 different options with a high F1 score of 0.93 in IMDB and 0.86 in Yelp. However, when DP-Prompt is employed, the sentiment F1 scores remain unchanged, while the author identification scores decrease by 46 and 25% in the case of IMDB, and 53 and 29% in the case of Yelp.

The text-level models are more accurate than the embedding-level models, with author identification scores of 0.99 (as opposed to 0.93) and 0.97 (as opposed to 0.86) in IMDB and Yelp, respectively, for clean reviews. When DP-Prompt is employed, the sentiment F1 scores remain unchanged, while the author identification scores decrease by 54 and 10% in the case of IMDB, and 73 and 24% in the case of Yelp. This illustrates that text-level attackers are more powerful.




Fig. 6.2 Comparison of DP-Prompt [841] (on ChatGPT) with various baselines. The first 4 plots show results for an attacker with embedding access, while the last 4 plots present results for an attacker with text access. Notably, it is evident that regardless of the chosen privacy level, DP-Prompt, when utilized with ChatGPT (GPT-3.5), exhibits significantly better utility compared to all baseline mechanisms. (a) IMDB (static). (b) IMDB (adaptive). (c) Yelp (static). (d) Yelp (adaptive). (e) IMDB (static). (f) IMDB (adaptive). (g) Yelp (static). (h) Yelp (adaptive)

Chapter 7

Memory-Efficient Fine-Tuning for Foundation Models



Abstract In the evolving landscape of foundation models, fine-tuning pre-trained models with first-order (FO) optimizers like SGD and Adam has become standard practice. However, as these models grow in size, the substantial memory overhead from back-propagation (BP) for FO gradient computation presents a significant challenge. Addressing this issue is crucial, particularly for applications like on-device training where memory efficiency is paramount. This chapter introduces a shift towards back-propagation-free, zeroth-order (ZO) optimization as a solution to reduce memory costs during model fine-tuning. We will explore both the algorithmic foundations of ZO optimization and its applications in memory-efficient fine-tuning.

7.1 Introduction

Fine-tuning pre-trained foundation models (FMs), such as large language models (LLMs), has become the *de facto* standard in modern ML paradigms [684, 726]. First-order (FO) optimizers, such as SGD [21] and Adam [404], have been the predominant choices for FM fine-tuning. However, as FMs continue to scale, they face significant memory overhead due to the back-propagation (BP) required for FO gradient computation. For instance, computing the gradient for the LLM OPT-13B requires 12 times more memory than model inference. This leads to the critical challenge of achieving *memory-efficient fine-tuning* for FMs. Addressing this challenge could also facilitate technological breakthroughs in related areas, such as on-device training, where memory efficiency is paramount [283, 1046].

To enhance memory efficiency, an emerging solution is to replace BP-required FO optimization methods with a BP-free optimizer during FM fine-tuning. This approach was initially proposed by Malladi et al. [568], where the FO gradient is approximated using a finite difference of function values. Despite its new application to FM fine-tuning, the underlying optimization principle is commonly known as *zeroth-order (ZO) optimization*, with the function value-based gradient estimate referred to as the ZO gradient estimate [200, 227, 254, 524, 606].

ZO optimization utilizes finite differences of function values to estimate FO gradients instead of relying on explicit gradient information. For a comprehensive overview of ZO optimization in both theory and practice, see [524]. This approach distinguishes itself from classical black-box optimization methods such as coordinate search [222], pattern search [815], evolutionary optimization [266], and Bayesian optimization [743]. Unlike these methods, ZO optimization requires minimal modifications to widely-used FO gradient-based algorithms, leveraging finite difference-based gradient estimation. ZO optimization also shares the provable convergence guarantees of FO methods [200, 606], and offers greater flexibility and scalability for handling larger problem sizes. Its potential has been recently unveiled in various deep learning applications, including adversarial attack and defense [124, 349, 832, 949, 1010, 1015], contrastive explanation [186], hyperparameter optimization [270, 871], generic policy design in reinforcement learning [445, 803], and resource-limited on-device training [271].

In this chapter, we will introduce key recent innovations in ZO optimization and its application to memory-efficient fine-tuning [110, 1009], with a focus on LLMs. Following [1009], we explore a broader range of ZO optimization methods beyond ZO-SGD, examining various tasks, model types, and evaluation metrics. This study aims to reveal the pros and cons of ZO optimization methods in terms of accuracy and efficiency. Building on these insights, further enhancements to ZO optimization-based LLM fine-tuning can be achieved using techniques such as block-wise descent, hybrid ZO and FO training, and gradient sparsity.

7.2 Algorithmic Foundations of ZO Optimization

In this section, we review the basics of ZO optimization and explain its rationale and application for fine-tuning FMs.

Basics of ZO Optimization ZO optimization serves as a gradient-free alternative to FO (first-order) optimization, approximating FO gradients through function value-based gradient estimates, which we call *ZO gradient estimates*. Thus, a ZO optimization method typically mirrors the algorithmic framework of its corresponding FO optimization counterpart. However, it substitutes the FO gradient with the ZO gradient estimate as the descent direction.

Various techniques exist for performing ZO gradient estimation [524]. The most representative method is the *randomized gradient estimator* (RGE) [200, 606], which relies on the finite difference of function values along a randomly chosen direction vector. RGE has also been used by Malladi et al. [568] to achieve memory-efficient fine-tuning for LLMs. Its preference in LLM fine-tuning is attributed to its *query efficiency*, i.e., a low number of function queries. To be specific, given a scalar-valued function $f(\mathbf{x})$ where $\mathbf{x} \in \mathbb{R}^d$ of dimension d , the RGE (denoted by $\hat{\nabla} f(\mathbf{x})$)

is expressed using central difference:

$$\hat{\nabla} f(\mathbf{x}) = \frac{1}{q} \sum_{i=1}^q \left[\frac{f(\mathbf{x} + \mu \mathbf{u}_i) - f(\mathbf{x} - \mu \mathbf{u}_i)}{2\mu} \mathbf{u}_i \right] \quad (\text{RGE})$$

where $\mathbf{u}_i$ is a random direction vector typically drawn from the standard Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$, q is the number of function queries, and $\mu > 0$ is a small perturbation stepsize (also known as smoothing parameter). Malladi et al. [568] employed Yet, it is worth noting that the query number q strikes a balance between the ZO gradient estimation variance and the query complexity. It has been shown in [200, 525] that the variance of RGE is roughly in the order of $O(d/q)$, where $O(\cdot)$ signifies the Big O notation. When reducing the query complexity is a primary task, (RGE) is typically specified using $q = 1$.

The rationale behind RGE stems from the concept of the *directional derivative* [200]: As $\mu \rightarrow 0$ (letting $q = 1$), the finite difference of function values in (RGE) approaches $f'(\mathbf{x}, \mathbf{u}) \doteq \mathbf{u}^T \nabla f(\mathbf{x})$, representing the directional derivative of $f(\mathbf{x})$ along the random direction $\mathbf{u} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Subsequently, RGE yields $\hat{\nabla} f(\mathbf{x}) \rightarrow f'(\mathbf{x}, \mathbf{u})\mathbf{u}$ as $\mu \rightarrow 0$. Moreover, the directional derivative provides us an unbiased gradient estimator of $\nabla f(\mathbf{x})$:

$$\mathbf{E}_{\mathbf{u}}[f'(\mathbf{x}, \mathbf{u})\mathbf{u}] = \mathbf{E}_{\mathbf{u}}[\mathbf{u}\mathbf{u}^T \nabla f(\mathbf{x})] = \nabla f(\mathbf{x}). \quad (7.1)$$

With the above background, the RGE $\hat{\nabla} f(\mathbf{x})$ can be interpreted as an approximation of the FO gradient $\nabla f(\mathbf{x})$ using the directional derivative.

As a byproduct of connecting RGE to (7.1), we obtain the directional derivative-based gradient estimate, $\nabla f(\mathbf{x}) \approx f'(\mathbf{x}, \mathbf{u})\mathbf{u}$, which is known as the forward gradient (**Forward-Grad**) [51, 701]. Different from RGE that relies solely on the finite difference of function values, Forward-Grad requires the use of forward mode automatic differentiation (AD) but eliminates the need for the backward evaluation in the implementation of deep model fine-tuning or training. In other words, Forward-Grad is BP-free and can serve as another alternative gradient estimation method that improves the memory efficiency of model fine-tuning.

Representative ZO Optimization Methods Next, we provide a brief overview of the ZO optimization methods in the literature. Specifically, we cover: **ZO-SGD** [254] that [568] has employed for LLM fine-tuning, ZO-SGD using sign-based gradient estimation (**ZO-SGD-Sign**) [523], ZO-SGD with momentum (**ZO-SGD-MMT**) [568], ZO-SGD with conservative gradient update (**ZO-SGD-Cons**), and the ZO variant of the Adam optimizer (**ZO-Adam**) [134].

The aforementioned methods can be unified into the following generic optimization framework in solving $\min_{\mathbf{x}} f(\mathbf{x})$:

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \eta_t h(\hat{\nabla} f(\mathbf{x}_t)), \quad (7.2)$$

where $\mathbf{x}_t$ denotes the updated solution at the t th iteration, $\eta_t > 0$ is the learning rate, and $h(\cdot)$ is a certain descent direction post-processing operation. In (7.2), we omit the inclusion of the stochastic mini-batch for empirical risk minimization for ease of presentation. For instance, ZO-SGD can be expressed as (7.2) when $h(\hat{\nabla} f(\mathbf{x})) = \hat{\nabla} f(\mathbf{x})$. Similarly, ZO-SGD-Sign can be derived if $h(\hat{\nabla} f(\mathbf{x})) = \text{sign}(\hat{\nabla} f(\mathbf{x}))$, where $\text{sign}(\cdot)$ represents element-wise sign operation. Another example is ZO-SGD-Cons by setting $h(\hat{\nabla} f(\mathbf{x})) = \arg \min_{\mathbf{g} \in \{0, -\hat{\nabla} f(\mathbf{x}), \hat{\nabla} f(\mathbf{x})\}} f(\mathbf{x}_t + \eta_t \mathbf{g})$.

The rationale for selecting the aforementioned ZO optimization approaches for FM fine-tuning is based on two key considerations: (1) We prioritize ZO optimization methods that require minimal modifications to the existing FO optimizer, ensuring ease of implementation for fine-tuning. (2) We focus on methods with distinct algorithmic characteristics, allowing us to explore a diverse range of optimization strategies for improving LLM performance. Regarding (2), we include ZO-SGD-Sign as it employs 1-bit gradient quantization and represents one of the simplest ZO optimization methods. Additionally, we include ZO-SGD-MMT and ZO-SGD-Cons as they incorporate certain forms of ‘adaptive learning’ into the descent step updates. The former utilizes momentum based on historical gradient information, while the latter allows for the heuristics-based selection of the descent direction. Furthermore, ZO-Adam is one of the most complex ZO optimizers due to its utilization of moving averages and adaptive learning rates.

Task Alignment in ZO Optimization for LLM Fine-Tuning Scaling up ZO optimization for deep model training, as discussed in [109], is exceedingly challenging due to its high variance, which is dependent on the model size. Nevertheless, model pre-training offers a unique advantage by enabling the fine-tuner to start from a well-optimized pre-trained model state. This graceful model initialization makes ZO optimization potentially scalable to FM fine-tuning tasks [568]. Even in this pretraining-finetuning paradigm, another crucial factor, which we call ‘**task alignment**’, still plays a key role in achieving satisfactory ZO fine-tuning performance. The ‘task alignment’ refers to aligning the fine-tuning task with the format of the pre-training task, given by e.g., the next token or sentence prediction for LLMs. It has shown in [247, 568] that downstream text classification tasks can be transformed into next token prediction tasks by introducing well-crafted input prompts. These prompts serve as bridges to align the fine-tuning tasks with the pre-training ones, facilitating ZO optimization when initiated from the pre-trained model.

As a warm-up experiment, Table 7.1 empirically justifies the importance of task alignment when applying ZO optimization to LLM fine-tuning on the simple binary classification task by comparing scenarios *with* and *without* the use of pre-defined prompts to achieve task alignment. We fine-tune the entire Roberta-Large [534] model on SST2 [767] and RTE [855] datasets with two selected ZO methods: ZO-SGD (i.e., MeZO in [568]) and ZO-Adam. We compare their performance with that of the FO method (FO-SGD). The task alignment is achieved with the template <CLS>SENTENCE. It was [terrible|great].<SEP> for SST dataset and another template <CLS>SENTENCE1? [Yes|No], SENTENCE2.<SEP> for RTE. As

Table 7.1 Test accuracy (%) of pretrained Roberta-Large model fine-tuned on SST2 and RTE tasks using ZO and FO optimization methods with (✓) and without (✗) text alignment. The evident performance degradation is highlighted in bold

Method	SST2			RTE		
	✓	✗	Difference	✓	✗	Difference
FO-SGD	91.6	91.5	0.1	70.9	61.4	9.5
ZO-SGD	89.4	79.2	10.2	68.7	60.4	8.3
ZO-Adam	89.8	79.2	10.6	69.2	58.7	10.5

we can see, without prompt-based text alignment, there is a substantial performance drop across ZO fine-tuning methods. Both ZO-SGD and ZO-Adam yield about 10% and 8% accuracy degradation on SST2 and RTE, respectively. In contrast, FO-SGD suffers less from the absence of task alignment. This suggests that the task alignment is particularly beneficial to ZO LLM fine-tuning. It is also worth noting that crafting effective prompts for task alignment can be non-trivial, as prompt design is context-dependent and can affect the fine-tuning performance.

7.3 Applying ZO Optimization for Memory-Efficient Fine-Tuning

In this section, we delve into the empirical performance of ZO optimization in fine-tuning LLMs. Our assessment includes evaluating both accuracy and efficiency across a range of downstream task complexities, from simple classification to more intricate reasoning tasks. Additionally, we consider various language model types to provide a comprehensive analysis of ZO optimization’s effectiveness.

LLM Fine-Tuning Tasks, Schemes, and Models We begin by introducing the tasks and the fine-tuning schemes. We focus on three tasks, considering their complexity from low to high, which include (1) the simplest binary classification task, Stanford Sentiment Treebank v2 (SST2) [767], (2) the question-answering task, Choice Of Plausible Alternatives (COPA) [706], and (3) the commonsense reasoning task, WinoGrande [721]. When evaluating memory efficiency, we also consider the task of multi-sentence reading comprehension (MultiRC) [393]. For LLM fine-tuning on these tasks, we explore four parameter-efficient fine-tuning (PEFT) schemes: full-tuning (FT) that fine-tunes the entire pre-trained model, low-rank adaptation (LoRA) by imposing low-rank weight perturbations [322], prefix-tuning (Prefix) by appending learnable parameters to token embedding [484], and prompt-tuning (Prompt) [528] by introducing a series of learnable tokens attached to the input to adapt the fixed model to downstream tasks. Furthermore, we incorporate several representative language models, including Roberta-Large [534], OPT [993], LLaMA2 [819], Vicuna [1022], and Mistral [372].

Table 7.2 Performance of LLM fine-tuning on SST2 over pretrained Roberta-Large and OPT/1.3B. Best performance among ZO methods (including Forward-Grad) is highlighted in **bold**

SST2	Roberta-Large				OPT-1.3B			
	FT	LoRA	Prefix	Prompt	FT	LoRA	Prefix	Prompt
FO-SGD	91.4	91.2	89.6	90.3	91.1	93.6	93.1	92.8
Forward-Grad	90.1	89.7	89.5	87.3	90.3	90.3	90.0	82.4
ZO-SGD	89.4	90.8	90.0	87.6	90.8	90.1	91.4	84.4
ZO-SGD-MMT	89.6	90.9	90.1	88.6	85.2	91.3	91.2	86.9
ZO-SGD-Cons	89.6	91.6	90.1	88.5	88.3	90.5	81.8	84.7
ZO-SGD-Sign	52.5	90.2	53.6	86.1	87.2	91.5	89.5	72.9
ZO-Adam	89.8	89.5	90.2	88.8	84.4	92.3	91.4	75.7

We evaluate ZO optimization-based fine-tuning using two sets of metrics: accuracy and efficiency. Accuracy measures the fine-tuned model’s test-time performance in specific tasks, such as test accuracy in classification tasks. Efficiency includes various measurements like memory efficiency (in terms of peak memory usage and GPU cost), query efficiency (i.e., the number of function queries required for ZO optimization), and run-time efficiency. These metrics collectively provide insights into the resources needed for ZO fine-tuning, helping assess its feasibility and cost-effectiveness in practical scenarios.

ZO Fine-Tuning on SST2: A Pilot Study In Table 7.2, we compare the performance of various BP-free and BP-based (FO-SGD) methods on the binary classification task using the SST2 dataset. We evaluate two model architectures: the medium-sized Roberta-Large and the larger OPT-1.3B. Key findings are summarized below. First, ZO-Adam [134] emerges as the most effective ZO method, achieving the best performance in 4 out of 8 fine-tuning settings. However, this comes at the cost of additional memory consumption, as ZO-Adam has the highest algorithmic complexity. Second, Forward-Grad [51, 701] proves to be competitive with ZO methods, especially in the FT (full-tuning) setting. This suggests that Forward-Grad may be suitable for larger-scale problems, making it a compelling baseline for ZO LLM fine-tuning. However, as the complexity of the fine-tuning scheme decreases (e.g., Prompt), the advantage of Forward-Grad over function value-based ZO methods diminishes. Third, The performance of ZO methods shows high variance, with fluctuating relative rankings across different scenarios despite extensive hyper-parameter tuning. For instance, the effectiveness of ZO-Adam drops significantly in the (OPT-1.3B, Prompt) setting. Moreover, the MeZO method (i.e., ZO-SGD) used in [568] is not always the top-performing ZO optimizer for LLM fine-tuning. This variance can be attributed to the high variance of the RGE [200, 606]. Fourth, ZO-SGD-Cons [396] and ZO-SGD-MMT [568] also demonstrate strong performance as ZO optimizers in LLM fine-tuning. However, ZO-SGD-Sign [523], the simplest ZO optimization method, tends to be the weakest approach except in the simplest fine-tuning setting (Prompt). These observations




Fig. 7.1 Results of OPT-13B on the tasks COPA and WinoGrande fine-tuned using ZO/FO optimizers in different PEFT settings

motivate us to extend our exploration of ZO methods across a broader spectrum of models and more complex tasks.

ZO Fine-Tuning on Downstream Tasks COPA and WinoGrande Under OPT-13B Extended from the experiments on SST2, Fig. 7.1 presents the fine-tuning performance on COPA and WinoGrande dataset using a larger model, OPT-13B. We summarize our key observations when the problem scales up and becomes more complicated. First, compared to previous results, the performance gap among different ZO methods has significantly widened. Additionally, the gap between FO and ZO methods has also increased. For example, in experiments with WinoGrande, FO methods (FO-SGD and FO-Adam) outperform all ZO methods by a large margin. This highlights the scalability bottleneck inherent to ZO methods when dealing with larger models and more complex tasks. Second, certain ZO methods exhibit exceptional stability across varied conditions. Despite a general trend towards variability, specific ZO methods like ZO-Adam and ZO-SGD-MMT demonstrate consistent performance. This stability is likely due to the integration of variance-reduced optimization techniques, such as momentum and adaptive learning rates, which make these algorithms more resilient to the variances of ZO gradient estimation [134]. Third, the LoRA tuning method consistently shows robustness when paired with various ZO algorithms. This resilience suggests that LoRA’s mechanism is inherently more adaptable to the variations in ZO optimization strategies, providing a stable and reliable tuning approach in diverse settings.

Table 7.3 The peak memory cost (in GB), the required GPU resources, and the runtime cost (in seconds) of each optimizer when fine-tuning the full OPT-13B model on MultiRC with an averaged 400 context length. The order of included optimizers is ranked based on the memory cost. The per-iteration runtime in seconds (s) is averaged over 100 iterations. Notably, Forward-Grad is marked by *, indicating its incompatibility with efficiency-enhancing techniques such as MP (mixed-precision training) and FP16 (half-precision training). Bold value represents best performance in each column

Optimizer	Memory ↓	Consumed GPUs ↓	Runtime cost
ZO-SGD	29 GB	1×A100	1.8 s
ZO-SGD-Cons	29 GB	1×A100	4.2 s
ZO-SGD-Sign	29 GB	1×A100	1.8 s
ZO-SGD-MMT	53 GB	1×A100	1.8 s
ZO-Adam	80 GB	2×A100	1.9 s
Forward-Grad*	138 GB	2×A100	19.8 s
FO-SGD	161 GB	3×A100	2.7 s
FO-Adam	257 GB	4×A100	2.8 s

Memory-Efficiency Analyses In Table 7.3, we present a comparison of the efficiency performance of various ZO and FO optimizers when fine-tuning the full OPT-13B model on the MultiRC dataset with a batch size of 4. We evaluate efficiency in terms of memory cost (in GB), GPU resource consumption (number of GPUs), and runtime cost per optimization iteration (in seconds). Key observations are as follows: First, almost all ZO methods (except ZO-Adam) demonstrate comparable levels of efficiency, requiring only a single GPU (A100) for LLM fine-tuning. This is expected as ZO methods involve relatively straightforward optimization steps, primarily based on function evaluations, as depicted in RGE. Among the examined ZO methods, ZO-Adam incurs higher memory costs due to its algorithmic complexity. Second, compared to FO methods, ZO optimization reduces runtime costs *per iteration*. For example, ZO-SGD reduces runtime by approximately 33.3% compared to FO-SGD. Third, Forward-Grad appears to be where ZO optimization methods lose their memory efficiency advantage over FO methods. Additionally, the substantial runtime cost of Forward-Grad compared to FO optimizers (FO-SGD and FO-Adam) is likely due to its incompatibility with mixed precision (MP) and FP16.

Furthermore, we examine the memory cost of LLM fine-tuning versus input sequence length. In Fig. 7.2, we compare the memory efficiency between ZO-SGD and FO-SGD across various sequence lengths (i.e., the token number per sample). We used synthetic texts generated from random sequences with specified shapes. The memory consumption of ZO-SGD remains consistent, whereas FO-SGD begins to demand significantly more memory as sequence length increases. This trend is particularly amplified for longer sequences (e.g., exceeding 700 tokens as depicted in Fig. 7.2), where the memory allocated for activations overwhelms that required for gradient storage.




Fig. 7.2 Memory comparison between FO-SGD and ZO-SGD full fine-tuning across various sequence lengths with a fixed effective batch size of 2. Memory evaluation was conducted using synthetic text generated from random sequences of the specified shapes. For shorter sequences (i.e., < 700), the memory usage of FO-SGD remains relatively stable since the memory consumption for storing gradients during BP surpasses that needed for activations




Fig. 7.3 Other FMs, e.g., InstructPix2Pix [79], Prompt Diffusion [878], and ControlNet [985], are customized using SD as the base model

Potential of ZO Fine-Tuning for Other Foundation Models Leveraging ZO optimization for memory-efficient fine-tuning can also be applied to other diffusion-based vision or vision-language models. Preliminary investigations of existing FMs in Fig. 7.3 support this potential. Notably, FMs for computer vision such as InstructPix2Pix [79], Prompt Diffusion [878], and ControlNet [985] are all

adaptations of the base model Stable Diffusion (SD) [709], making them suitable candidates for model fine-tuning via ZO optimization. One can explore the fine-tuning performance, such as the accuracy and convergence rate of ZO optimization, and analyze its interrelationship with model pre-training and adaptation. Additionally, examining the loss landscape of both the pre-trained and fine-tuned models using ZO optimization can provide insights into the potential impact of pre-training and ZO optimization on memory efficiency.

Chapter 8

Large Language Models Meet Time Series



Abstract Time series forecasting holds significant importance in many real-world dynamic systems and has been extensively studied. Unlike natural language process (NLP) and computer vision (CV), where a single large model can tackle multiple tasks, models for time series forecasting are often specialized, necessitating distinct designs for different tasks and applications. While pre-trained foundation models have made impressive strides in NLP and CV, their development in time series domains has been constrained by data sparsity. Recent studies have revealed that large language models (LLMs) possess robust pattern recognition and reasoning abilities over complex sequences of tokens. However, the challenge remains in effectively aligning the modalities of time series data and natural language to leverage these capabilities. In this chapter, we present TIME-LLM (Jin et al., Time-LLM: Time series forecasting by reprogramming large language models. In: The Twelfth International Conference on Learning Representations (2024)), a reprogramming framework to repurpose pretrained LLMs for general time series forecasting with the backbone language models kept intact. TIME-LLM reprograms the input time series with text prototypes before feeding it into the frozen LLM to align the two modalities. To augment the LLM's ability to reason with time series data, Prompt-as-Prefix (PaP) is proposed to enrich the input context and direct the transformation of reprogrammed input patches. The transformed time series patches from the LLM are finally projected to obtain the forecasts. The comprehensive evaluations demonstrate that TIME-LLM is a powerful time series learner that outperforms state-of-the-art, specialized forecasting models.

8.1 Introduction

Time series forecasting is a critical capability across many real-world dynamic systems, with applications ranging from demand planning [451] and inventory optimization [477] to energy load forecasting [514] and climate modeling [731]. Each time series forecasting task typically requires extensive domain expertise and task-specific model designs. This stands in stark contrast to foundation language

models like GPT-3 [81], GPT-4 [621], Llama [819], *inter alia*, which can perform well on a diverse range of NLP tasks in a few-shot or even zero-shot setting.

Pre-trained foundation models, such as large language models (LLMs), have driven rapid progress in computer vision (CV) and natural language processing (NLP). While time series modeling has not benefited from the same significant breakthroughs, LLMs’ impressive capabilities have inspired their application to time series forecasting. Several desiderata exist for leveraging LLMs to advance forecasting techniques: **Generalizability**. LLMs have demonstrated a remarkable capability for few-shot and zero-shot transfer learning [81]. This suggests their potential for generalizable forecasting across domains without requiring per-task retraining from scratch. In contrast, current forecasting methods are often rigidly specialized by domain. **Data efficiency**. By leveraging pre-trained knowledge, LLMs have shown the ability to perform new tasks with only a few examples. This data efficiency could enable forecasting for settings where historical data is limited. In contrast, current methods typically require abundant in-domain data. **Reasoning**. LLMs exhibit sophisticated reasoning and pattern recognition capabilities [154, 588, 872]. Harnessing these skills could allow making highly precise forecasts by leveraging learned higher-level concepts. Existing non-LLM methods are largely statistical without much innate reasoning. **Multimodal knowledge**. As LLM architectures and training techniques improve, they gain more diverse knowledge across modalities like vision, speech, and text [561]. Tapping into this knowledge could enable synergistic forecasting that fuses different data types. Conventional tools lack ways to jointly leverage multiple knowledge bases. **Easy optimization**. LLMs are trained once on massive computing and then can be applied to forecasting tasks without learning from scratch. Optimizing existing forecasting models often requires significant architecture search and hyperparameter tuning [1041]. In summary, LLMs offer a promising path to make time series forecasting more general, efficient, synergistic, and accessible compared to current specialized modeling paradigms. Thus, adapting these powerful models for time series data can unlock significant untapped potential.

The realization of the above benefits hinges on the effective alignment of the modalities of time series data and natural language. However, this is a challenging task as LLMs operate on discrete tokens, while time series data is inherently continuous. Furthermore, the knowledge and reasoning capabilities to interpret time series patterns are not naturally present within LLMs’ pre-training. Therefore, it remains an open challenge to unlock the knowledge within LLMs and activate their ability for general time series forecasting in a way that is accurate, data-efficient, and task-agnostic.

In this work, we present TIME-LLM [379], a reprogramming framework to adapt LLMs for time series forecasting while keeping the backbone model intact. The core idea is to *reprogram* the input time series into text prototype representations that are more naturally suited to language models’ capabilities. To further augment the model’s reasoning about time series concepts, we introduce *Prompt-as-Prefix* (PaP), a novel idea in enriching the input time series with additional context and providing task instructions in the modality of natural language. This provides declarative

guidance about desired transformations to apply to the reprogrammed input. The output of the language model is then projected to generate time series forecasts. Our comprehensive evaluation demonstrates that LLMs can act as effective few-shot and zero-shot time series learners when adopted through this reprogramming approach, outperforming specialized forecasting models. By leveraging LLMs’ reasoning capability while keeping the models intact, TIME-LLM points the way toward multimodal foundation models that can excel on both language and sequential data tasks. The proposed reprogramming framework offers an extensible paradigm for imbuing large models with new capabilities beyond their original pre-training.

8.2 Background and Related Work

Task-Specific Learning Most time series forecasting models are crafted for specific tasks and domains (e.g., traffic prediction), and trained end-to-end on small-scale data. An illustration is in Fig. 8.1a. For example, ARIMA models are designed for univariate time series forecasting [75], LSTM networks are tailored for sequence modeling [305], and temporal convolutional networks [39] and transformers [895] are developed for handling longer temporal dependencies. While achieving good performance on narrow tasks, these models lack versatility and generalizability to diverse time series data.

In-Modality Adaptation Relevant research in CV and NLP has demonstrated the effectiveness of pre-trained models that can be fine-tuned for various downstream tasks without the need for costly training from scratch [81, 182, 819]. Inspired by these successes, recent studies have focused on the development of time series pre-trained models (TSPTMs). The first step among them involves time series pre-training using different strategies like supervised [217] or self-supervised learning [177, 982, 1005]. This allows the model to learn representing various input time series. Once pre-trained, it can be fine-tuned on similar domains to learn how to perform specific tasks [801]. An example is in Fig. 8.1b. The development of




Fig. 8.1 Schematic illustration of reprogramming large language models (LLMs) in comparison of (a) task-specific learning and (b) model fine-tuning. Our proposal investigates and demonstrates (c) how to effectively reprogram open-sourced LLMs as powerful time series learners where well-developed time series pre-trained models are not readily available

TSPTMs leverages the success of pre-training and fine-tuning in NLP and CV but remains limited on smaller scales due to data sparsity.

Cross-Modality Adaptation Building on in-modality adaptation, recent work has further explored transferring knowledge from powerful pre-trained foundations models in NLP and CV to time series modeling, through techniques such as multimodal fine-tuning [953] and model reprogramming [120]. Our approach aligns with this category; however, there is limited pertinent research available on time series. An example is Voice2Series [936], which adapts an acoustic model (AM) from speech recognition to time series classification. It transforms a time series into a format suitable for the AM and map the output to new labels, allowing for leveraging the representation learning ability of AMs trained on massive voice datasets for quick adaptation on time series. Recently, [103] proposes LLM4TS for time series forecasting using LLMs. It designs a two-stage fine-tuning process on the LLM—first supervised pre-training on time series, then task-specific fine-tuning. [1040] leverages pre-trained language models without altering their self-attention and feedforward layers. This model is fine-tuned and evaluated on various time series analysis tasks and demonstrates comparable or state-of-the-art performance by transferring knowledge from natural language pre-training. Distinct from these approach, we neither edit the input time series directly nor fine-tune the backbone LLM. Instead, as illustrated in Fig. 8.1c, we propose reprogramming time series with the source data modality along with prompting to unleash the potential of LLMs as effective time series machines.

8.3 Time-LLM

The model architecture of TIME-LLM is depicted in Fig. 8.2. We focus on reprogramming an embedding-visible language foundation model, such as Llama [819] and GPT-2 [680], for general time series forecasting *without* requiring any fine-tuning of the backbone model. Specifically, we consider the following problem: given a sequence of historical observations $\mathbf{X} \in \mathbb{R}^{N \times T}$ consisting of N different 1-dimensional variables across T time steps, we aim to reprogram a large language model $f(\cdot)$ to understand the input time series and accurately forecast the readings at H future time steps, denoted by $\hat{\mathbf{Y}} \in \mathbb{R}^{N \times H}$, with the overall objective to minimize the mean square errors between the ground truths $\mathbf{Y}$ and predictions, i.e., $\frac{1}{H} \sum_{h=1}^H \|\hat{\mathbf{Y}}_h - \mathbf{Y}_h\|_F^2$.

TIME-LLM encompasses three main components: (1) input transformation, (2) a pre-trained and frozen LLM, and (3) output projection. Initially, a multivariate time series is partitioned into N univariate time series, which are subsequently processed independently [611]. The i -th series is denoted as $\mathbf{X}^{(i)} \in \mathbb{R}^{1 \times T}$, which undergoes normalization, patching, and embedding prior to being reprogrammed with learned text prototypes to align the source and target modalities. Then, we augment the LLM’s time series reasoning ability by prompting it together with reprogrammed




Fig. 8.2 The model framework of TIME-LLM [379]. Given an input time series, we first tokenize and embed it via (i) patching along with a (ii) customized embedding layer. (iii) These patch embeddings are then reprogrammed with condensed text prototypes to align two modalities. To augment the LLM’s reasoning ability, (iv) additional prompt prefixes are added to the input to direct the transformation of input patches. (v) The output patches from the LLM are projected to generate the forecasts

patches to generate output representations, which are projected to the final forecasts $\hat{\mathbf{Y}}^{(i)} \in \mathbb{R}^{1 \times H}$.

We note that only the parameters of the lightweight input transformation and output projection are updated, while the backbone language model is frozen. In contrast to vision-language and other multimodal language models, which usually fine-tune with paired cross-modality data, TIME-LLM is directly optimized and becomes readily available with only a small set of time series and a few training epochs, maintaining high efficiency and imposing fewer resource constraints compared to building large domain-specific models from scratch or fine-tuning them. To further reduce memory footprints, various off-the-shelf techniques (e.g., quantization) can be seamlessly integrated for slimming TIME-LLM.

8.3.1 Model Structure

Input Embedding Each input channel $\mathbf{X}^{(i)}$ is first individually normalized to have zero mean and unit standard deviation via reversible instance normalization (RevIN) in mitigating the time series distribution shift [402]. Then, we divide $\mathbf{X}^{(i)}$ into several consecutive overlapped or non-overlapped patches [611] with length L_p ;

thus the total number of input patches is $P = \lfloor \frac{(T-L_p)}{S} \rfloor + 2$, where S denotes the horizontal sliding stride. The underlying motivations are two-fold: (1) better preserving local semantic information by aggregating local information into each patch and (2) serving as tokenization to form a compact sequence of input tokens, reducing computational burdens. Given these patches $\mathbf{X}_p^{(i)} \in \mathbb{R}^{P \times L_p}$, we embed them as $\hat{\mathbf{X}}_p^{(i)} \in \mathbb{R}^{P \times d_m}$, adopting a simple linear layer as the patch embedder to create dimensions d_m .

Patch Reprogramming Here we reprogram patch embeddings into the source data representation space to align the modalities of time series and natural language to activate the backbone’s time series understanding and reasoning capabilities. A common practice is learning a form of “noise” that, when applied to target input samples, allows the pre-trained source model to produce the desired target outputs without requiring parameter updates. This is technically feasible for bridging data modalities that are identical or similar. Examples include repurposing a vision model to work with cross-domain images [590] or reprogramming an acoustic model to handle time series data [936]. In both cases, there are explicit, learnable transformations between the source and target data, allowing for the direct editing of input samples. However, time series can neither be directly edited nor described losslessly in natural language, posing significant challenges to directly bootstrap the LLM for understanding time series without resource-intensive fine-tuning.

To close this gap, we propose reprogramming $\hat{\mathbf{X}}_p^{(i)}$ using pre-trained word embeddings $\mathbf{E} \in \mathbb{R}^{V \times D}$ in the backbone, where V is the vocabulary size. Nevertheless, there is no prior knowledge indicating which source tokens are directly relevant. Thus, simply leveraging $\mathbf{E}$ will result in large and potentially dense reprogramming space. A simple solution is to maintain a small collection of text prototypes by linearly probing $\mathbf{E}$, denoted as $\mathbf{E}' \in \mathbb{R}^{V' \times D}$, where $V' \ll V$. An illustration is in Fig. 8.3a. Text prototypes learn connecting language cues, e.g., “long steady” (blue lines) and “short up” (red lines), which are then combined to represent the local patch information (e.g., “short up then down steadily” for characterizing patch 5) without leaving the space where the language model is pre-trained. This approach is




Fig. 8.3 Illustration of (a) patch reprogramming and (b) Patch-as-prefix versus prompt-as-prefix

efficient and allows for the adaptive selection of relevant source information. To realize this, we employ a multi-head cross-attention layer. Specifically, for each head $k = \{1, \dots, K\}$, we define query matrices $\mathbf{Q}_k^{(i)} = \hat{\mathbf{X}}_P^{(i)} \mathbf{W}_k^Q$, key matrices $\mathbf{K}_k^{(i)} = \mathbf{E}' \mathbf{W}_k^K$, and value matrices $\mathbf{V}_k^{(i)} = \mathbf{E}' \mathbf{W}_k^V$, where $\mathbf{W}_k^Q \in \mathbb{R}^{d_m \times d}$ and $\mathbf{W}_k^K, \mathbf{W}_k^V \in \mathbb{R}^{D \times d}$. Specifically, D is the hidden dimension of the backbone model, and $d = \lfloor \frac{d_m}{K} \rfloor$. Then, we have the operation to reprogram time series patches in each attention head defined as:

$$\mathbf{Z}_k^{(i)} = \text{ATTENTION}(\mathbf{Q}_k^{(i)}, \mathbf{K}_k^{(i)}, \mathbf{V}_k^{(i)}) = \text{SOFTMAX}\left(\frac{\mathbf{Q}_k^{(i)} \mathbf{K}_k^{(i)\top}}{\sqrt{d_k}}\right) \mathbf{V}_k^{(i)}. \quad (8.1)$$

By aggregating each $\mathbf{Z}_k^{(i)} \in \mathbb{R}^{P \times d}$ in every head, we obtain $\mathbf{Z}^{(i)} \in \mathbb{R}^{P \times d_m}$. This is then linearly projected to align the hidden dimensions with the backbone model, yielding $\mathbf{O}^{(i)} \in \mathbb{R}^{P \times D}$.

Prompt-as-Prefix Prompting serves as a straightforward yet effective approach task-specific activation of LLMs [953]. However, the direct translation of time series into natural language presents considerable challenges, hindering both the creation of instruction-following datasets and the effective utilization of on-the-fly prompting without performance compromise [932]. Recent advancements indicate that other data modalities, such as images, can be seamlessly integrated as the prefixes of prompts, thereby facilitating effective reasoning based on these inputs [831]. Motivated by these findings, and to render our approach directly applicable to real-world time series, we pose an alternative question: *can prompts act as prefixes to enrich the input context and guide the transformation of reprogrammed time series patches?* We term this concept as *Prompt-as-Prefix* (PaP) and observe that it significantly enhances the LLM’s adaptability to downstream tasks while complementing patch reprogramming.

An illustration of the two prompting approaches is in Fig. 8.3b. In *Patch-as-Prefix*, a language model is prompted to predict subsequent values in a time series, articulated in natural language. This approach encounters certain constraints: (1) language models typically exhibit reduced sensitivity in processing high-precision numerals without the aid of external tools, thereby presenting substantial challenges in accurately addressing practical forecasting tasks over long horizons; (2) intricate, customized post-processing is required for different language models, given that they are pre-trained on diverse corpora and may employ different tokenization types in generating high-precision numerals with precision and efficiency. This results in forecasts being represented in disparate natural language formats, such as ['0', '.', '6', '1'] and ['0', '.', '61'], to denote the decimal 0.61.

Prompt-as-Prefix, on the other hand, tactfully avoids these constraints. In practice, we identify three pivotal components for constructing effective prompts: (1) dataset context, (2) task instruction, and (3) input statistics. A prompt example is in Fig. 8.4. The dataset context furnishes the LLM with essential background information concerning the input time series, which often exhibits distinct charac-

The Electricity Transformer Temperature (ETT) indicates the electric power long-term deployment. Each data point consists of the target oil temperature and 6 power load features ...
Below is the information about the input time series:

[BEGIN DATA]

[Domain]: We usually observe that electricity consumption peaks at noon, with a significant increase in transformer load

[Instruction]: Predict the next $\langle H \rangle$ steps given the previous $\langle T \rangle$ steps information attached

[Statistics]: The input has a minimum of $\langle \text{min\_val} \rangle$, a maximum of $\langle \text{max\_val} \rangle$, and a median of $\langle \text{median\_val} \rangle$. The overall trend is $\langle \text{upward or downward} \rangle$. The top five lags are $\langle \text{lag\_val} \rangle$.

[END DATA]

Fig. 8.4 Prompt example. $\langle \rangle$ and $\langle \rangle$ are task-specific configurations and calculated input statistics

teristics across various domains. Task instruction serves as a crucial guide for the LLM in the transformation of patch embeddings for specific tasks. We also enrich the input time series with additional crucial statistics, such as trends and lags, to facilitate pattern recognition and reasoning.

Output Projection Upon packing and feedforwarding the prompt and patch embeddings $\mathbf{O}^{(i)}$ through the frozen LLM as shown in Fig. 8.2, we discard the prefixal part and obtain the output representations. Following this, we flatten and linear project them to derive the final forecasts $\hat{\mathbf{Y}}^{(i)}$.

8.4 Performance Evaluation

Experiment Setup To ensure a fair comparison, we adhere to the experimental configurations in [906] across all baselines with a unified evaluation pipeline.¹ We use Llama-7B [819] as the default backbone unless stated otherwise.

Baselines We compare with the SOTA time series models, and we cite their performance from [1040] if applicable. Our baselines include a series of

¹ <https://github.com/thuml/Time-Series-Library>.

Transformer-based methods: PatchTST [611], ESTformer [904], Non-Stationary Transformer [535], FEDformer [1039], Autoformer [907], Informer [1034], and Reformer [412]. We also select a set of recent competitive models, including GPT4TS [1040], LLMTime [269], DLinear [974], TimesNet [906], and LightTS [999]. In short-term forecasting, we further compare TIME-LLM with N-HiTS [101] and N-BEATS [624].

Numerical Results of Long-Term Forecasting We evaluate on ETTh1, ETTh2, ETTm1, ETTm2, Weather, Electricity (ECL), Traffic, and ILI, which have been extensively adopted for benchmarking long-term forecasting models [906]. The input time series length T is set as 512, and we use four different prediction horizons $H \in \{96, 192, 336, 720\}$. The evaluation metrics include mean square error (MSE) and mean absolute error (MAE). The results are shown in Table 8.1, where TIME-LLM outperforms all baselines in most cases and significantly so to the majority of them. The comparison with GPT4TS [1040] is particularly noteworthy. GPT4TS is a very recent work that involves fine-tuning on the backbone language model. We note average performance gains of **12** and **20%** over GPT4TS and TimesNet, respectively. When compared with the SOTA task-specific Transformer model PatchTST, by reprogramming the smallest Llama, TIME-LLM realizes an average MSE reduction of 1.4%. Relative to the other models, e.g., DLinear, our improvements are also pronounced, exceeding **12%**.

Numerical Results of Short-Term Forecasting We choose the M4 benchmark [567] as the testbed, which contains a collection of marketing data in different sampling frequencies. The prediction horizons in this case are relatively small and in [6, 48]. The input lengths are twice as prediction horizons. The evaluation metrics are symmetric mean absolute percentage error (SMAPE), mean absolute scaled error (MSAE), and overall weighted average (OWA). The results with unified seeds across all methods are in Table 8.2. TIME-LLM consistently surpasses all baselines, outperforming GPT4TS by **8.7%**. TIME-LLM remains competitive even when compared with the SOTA model, N-HiTS [101], w.r.t. MASE and OWA.

In [379], the authors also studied the few-shot learning and zero-shot forecasting settings. TIME-LLM remarkably excels over all baseline methods, and the authors attribute this to the successful knowledge activation in the reprogrammed LLM.

Table 8.2 Short-term time series forecasting results on M4. The forecasting horizons are in [6, 48] and the three rows provided are weighted averaged from all datasets under different sampling intervals. A lower value indicates better performance. **Red**: the best, **Blue**: the second best

Methods	TIME-LLM	GPT4TS	TimesNet	PatchTST	N-HiTS	N-BEATS	ETSformer	LightTS	DL-linear	FEDformer	Stationary	Autoformer	Informr	Reformer
Average		[379]	[1040]	[906]	[611]	[101]	[624]	[904]	[999]	[974]	[1039]	[535]	[907]	[412]
	SMAPE	11.983	12.69	12.88	12.059	12.035	12.25	14.718	13.525	13.639	13.16	12.780	12.909	18.200
	MASE	1.595	1.808	1.836	1.623	1.625	1.698	2.408	2.111	2.095	1.775	1.756	1.771	4.223
	OWA	0.859	0.94	0.955	0.869	0.869	0.896	1.172	1.051	1.051	0.949	0.930	0.939	1.775

Chapter 9

Large Language Models Meet Speech Recognition



Abstract Advancements in deep neural networks have allowed automatic speech recognition (ASR) systems to attain human parity on several publicly available clean speech datasets. However, even state-of-the-art ASR systems experience performance degradation when confronted with adverse conditions, as a well-trained acoustic model is sensitive to variations in the speech domain, e.g., background noise. On the other hand, linguistic knowledge in large language models (LLMs) can be used to infer the meaning of ambiguous spoken terms from contextual cues, thereby reducing the dependency on the auditory system. Based on rich linguistic knowledge and powerful reasoning ability of LLMs, this chapter presents recent studies in using LLMs for generative error correction (GER) in ASR to improve recognition results.

9.1 Introduction

Automatic speech recognition (ASR) is an essential technology that enables efficient and accurate transcription of spoken languages. This capability facilitates access to information and enhances communication across various domains, including education [86], healthcare [437], and business [297]. Driven by the recent advances in deep learning, remarkable success has been achieved on several ASR tasks through end-to-end training techniques [102, 137, 191, 263, 264, 274, 934]. However, a major challenge of applying ASR in practical conditions lies in effectively handling variations in speech caused by different factors such as background noise [112], speaker accent [836], and speaking styles [14, 797]. These adverse factors are common and inevitable in speech signal, significantly affecting the accuracy of the recognition results [472].

Humans demonstrate remarkable robustness when faced with the above variations in acoustic environment, as the human recognition system does not only rely on acoustic cues—we usually speculate the ambiguous or distorted spoken terms based on speech context and our inherent linguistic knowledge. Similarly, current ASR system typically employs an independent language model (LM) for rescoring during the decoding process [240, 334, 384, 816]. As shown in




Fig. 9.1 The left part shows the pipeline to generate the N-best hypotheses using a vanilla ASR engine with beam search decoding. The right part counts the probabilities of case (i) and case (ii) on the test set of LibriSpeech dataset. It indicates the discarded information in 2nd ~ 20th utterances. Green and red T_i in “Exp” respectively denote correct and wrong tokens compared with ground-truth

Fig. 9.1, given N-best hypotheses generated by an ASR engine with beam search decoding, a trained language model (LM) can be used to re-score each utterance and select the one with the highest likelihood (referred to as the 1st utterance) as the output of the ASR; whereas, the other sentences (the 2nd– N th utterances) are discarded. However, it is widely believed [648] that the N-best list contains useful information [304, 475, 842], as each hypothesis is an independent textual representation of the input speech. Consequently, discarded sentences might also carry correct tokens for accurately predicting the true transcription. To validate this belief, we have conducted experiments on the LibriSpeech dataset [637], counting the probabilities of two scenarios observed during LM rescoring: (i) the discarded utterances contain a better candidate with lower word error rate (WER), and (ii) the other discarded hypotheses can provide the right answer for the wrong tokens in 1st utterance. The statistical results of 2nd ~ 20th utterances are shown in the left part of Fig. 9.1. Taking 2nd discarded utterance as example, it has a 14% probability of having a lower WER than the 1st utterance. Furthermore, given a wrong token in 1st utterance, there is a 34% probability of finding the correct token in the 2nd utterance.

To better mine the information in N-best hypotheses, the authors in [113] propose the first attempt on publicly available ASR generative error correction benchmark that directly predicts a true transcription, rather than selecting a candidate from the N-best list. This benchmark is named *HyParadise (HP)*, which comprises various open source N-best hypotheses provided by state-of-the-art ASR systems and their paired true transcriptions. Considering real-life applications, HP dataset covers various challenging speech domains, including scenarios with background noise, specific contexts, and speaker accents. Different evaluation settings are defined to mimic the deployment of ASR systems in real-world scenarios:

1. *Zero-shot Learning*. In this setting, only test set hypotheses are available for inference. This corresponds to applying a well-trained ASR model to new scenarios without any training data.

2. *Few-shot Learning*. A few in-domain hypotheses with true transcription are available for training. This setting aims to address domain-specific ASR tasks with a few manual annotations.
3. *Fine-tuning*. A sufficient training set is available to learn the mapping between hypotheses and transcription.

Further, in [330], the authors extend the GER benchmark [113] to noisy conditions, as well as propose a Robust HyParadise (RobustHP) dataset with 113K hypotheses-transcription pairs from various ASR corpus in common noisy scenarios. The authors also propose new methods to teach LLMs to denoise the N-best hypotheses for GER. Specifically, they propose to extract a noise embedding in *language space* to represent the noise conditions of source speech, by measuring the diversity of N-best hypotheses list from ASR decoding. The insight behind is that, the worse noisy conditions (more challenging noise type or lower signal-to-noise ratio), the higher uncertainty of ASR beam search decoding, and thus results in more diverse N-best hypotheses. Extracted from the language space of hypotheses instead of audio space, the noise embedding can be well incorporated into LLM tuning to improve GER, which can be viewed as a novel *language-space denoising* process. Moreover, in order to enhance its representation ability of audio noise, they design a knowledge distillation (KD) approach via mutual information estimation [56] to distill the real noise information in audio embeddings to the extracted language embedding.

9.2 Background and Related Work

9.2.1 ASR Rescoring and Error Correction

In order to improve the linguistic acceptability of ASR results, language model (LM) rescoring has been widely employed and achieved stable performance gain for ASR systems [32, 586, 756]. Typically, an external LM is trained separately and utilized to re-score the N-best list of hypotheses generated by ASR decoding with beam search. Various approaches for LM integration have been proposed, such as shallow fusion [144, 384, 816, 975], deliberation [286, 323–325, 868, 916], component fusion [745], and cold fusion [781]. Some authors have used pre-trained LM models to replace trainable LMs [722, 839], and the log-likelihood of each hypothesis is computed using unidirectional models, e.g., GPT-2, or pseudo-log-likelihood using bidirectional models like BERT [182] and RoBERTa [534]. In ASR, LMs are also widely used for the error correction task in different languages [265, 902], leveraging only the 1-best hypothesis generated by the ASR model [201, 450, 569, 748, 992, 1018]. Furthermore, more recent works [448, 449, 557] utilize a candidates list after decoding for error correction. Though Grammatical Error Correction (GEC) has been actively explored [168, 875, 934], ASR error correction is distinct

with GER due to the arbitrariness of the spoken language [14], which requires the efforts from both speech and natural language processing communities [152].

9.2.2 Noise-Robust ASR

Neural ASR has achieved human-level performance in recent years but its noise-robustness in the real world remains a challenge [421]. Recent noise-robust ASR methods make some progress by mapping noisy speech features to clean space (i.e., denoise) before recognition [472, 638]. For instance, speech enhancement serves as a denoising front-end [236, 329] to improve speech quality for ASR [331–333], domain adversarial training aims to learn noise-invariant speech features [661], and the recently popular ASR foundation model proposes to use web-scale data and various preprocessing steps for denoising [677].

9.2.3 HyPoradise (HP) Benchmarks

In [113], the authors employ two state-of-the-art ASR models, namely WavLM [128] and Whisper [676] for N -best hypotheses generation. To cover common scenarios of ASR task, e.g., noisy background and speaker accent, the authors selected 9 representative corpora to compose the HP dataset. In total, they collected more than 334K pairs of hypotheses list and transcription to form the HP dataset, including training and test sets.

In [330], the authors extend the HyPoradise benchmark [113] to noise-robust ASR. Given an input noisy speech X_n , the pre-trained ASR model first transcribes it into N -best hypotheses $\mathcal{Y}_N = \{Y_1, Y_2, \dots, Y_N\}$ by beam search decoding. The goal of GER is to learn a hypotheses-to-transcription (H2T) mapping $\mathcal{M}_{\text{H2T}}$ that predicts the transcription Y based on N -best hypotheses list $\mathcal{Y}_N$:

$$Y = \mathcal{M}_{\text{H2T}}(\mathcal{Y}_N), \quad (9.1)$$

Given the ground-truth transcription Y^* , we can finetune the LLM to learn $\mathcal{M}_{\text{H2T}}$ in an auto-regressive manner, where the cross-entropy loss $\mathcal{L}_{\text{H2T}}$ is formulated as:

$$\mathcal{L}_{\text{H2T}} = \sum_{t=1}^T -\log \mathcal{P}_{\theta}(y_t^* | y_{t-1}^*, \dots, y_1^*, \mathcal{Y}_N), \quad (9.2)$$

where y_t^* is the t -th token of Y^* , and θ denotes the learnable parameters in LLM (e.g., adapter).

Correspondingly, they develop a Robust HyPoradise dataset by collecting hypotheses-transcription (HT) pairs from common noisy ASR corpus. They employ

Whisper Large-V2 [677] to transcribe the noisy speech into N-best hypotheses (N is set to 5), resulting in 113K HT pairs in total from various noise domains.

9.3 Noise-Aware Generative Error Correction

In this section, we present the noise-aware generative error correction (RobustGER) approach proposed in [330]. The left part of Fig. 9.2 presents the overall framework of RobustGER. First, the noisy speech X_n is sent into a pre-trained ASR model to generate N-best hypotheses $\mathcal{Y}_N = \{Y_1, Y_2, \dots, Y_N\}$, where $N = 5$. Following that, it extracts a language-space noise embedding E_{LN} from the N-best list $\mathcal{Y}_N$ to represent the noise conditions of source speech X_n . As depicted in the right part of Fig. 9.2, such noise embedding measures the diversity of N-best hypotheses on both utterance and token levels, which perceives the noise information in input speech.

Furthermore, to enhance the noise representation ability, a knowledge distillation (KD) approach is used to distill the real noise information in source speech X_n to the extracted language-space noise embedding E_{LN} . Specifically, we employ the audio embedding $\mathcal{E}_{ASR}(X_n)$ from ASR encoder for distillation. Finally, we add an instruction onto the N-best hypotheses and sent them into LLM to predict the true transcription (i.e., GER), with the language embedding incorporated for denoising. Specifically, we add a minus sign before the noise embedding E_{LN} to indicate “denoise”. Such minus embedding is then sent to teach LLM to do language-space




Fig. 9.2 Left: the RobustGER framework [330] that leverages efficient LLM finetuning to learn mapping from ASR N-best hypotheses to ground-truth transcription, where we propose a language-space noise embedding with audio distillation to denoise GER process. **Right:** the extraction of language-space noise embedding from N-best hypotheses by measuring its diversity, where we calculate the utterance- and token-level embedding differences between each pair of hypotheses in the N-best list

denoising. Therefore, Eq.(9.1) should be re-written as:

$$Y = \mathcal{M}_{\text{H2T}}(\mathcal{Y}_N; -E_{\text{LN}}), \quad (9.3)$$

The $\mathcal{M}_{\text{H2T}}$ denotes H2T mapping by efficient LLM finetuning, where we follow the adapter tuning from previous works [935, 989]. Similar to Eq.(9.2), we follow the original GER benchmark for optimization.

9.3.1 Language-Space Noise Embedding

As directly incorporating audio-space noise embedding into LLM finetuning could harm its stability and performance [246, 989], we propose an alternative to extract language-space noise embedding from N-best hypotheses to represent the noise conditions of source speech. The key idea is to perceive the audio noise from the diversity of N-best hypotheses, i.e., the worse noisy conditions (more challenging noise type or lower SNR), the higher uncertainty of ASR beam search decoding, and thus results in more diverse N-best hypotheses.

As illustrated in the right part of Fig. 9.2, we extract the noise embedding on both utterance and token levels to capture rich diversity information: (1) *Utterance-level*: examine the diversity inside N-best list in terms of the entire utterance’s semantic meaning, which indicates the affect of audio noise on the global semantics of hypotheses; (2) *Token-level*: examine the distribution of N-best hypothesis in terms of all the tokens inside, which is similar to edit distance and thus directly corresponds to the WER metric. These two embeddings are finally combined to form the resulted noise embedding, i.e., $E_{\text{LN}} = [E_{\text{LN}}^{\text{utt}}; E_{\text{LN}}^{\text{tok}}]$. Specifically, we employ sentence-BERT (SBERT) [699] to obtain the embeddings from raw text, which contains rich language-space semantic information.

Utterance-Level Noise Embedding Given N-best hypotheses $\mathcal{Y}_N = \{Y_1, Y_2, \dots, Y_N\}$, we first obtain their sentence embeddings by SBERT encoder $\mathcal{E}_{\text{sbert}}$ and then calculate their diversity as:

$$E_{\text{LN}}^{\text{utt}} = \text{Concat}\{[\mathcal{E}_{\text{sbert}}(Y_i) - \mathcal{E}_{\text{sbert}}(Y_j)]_{i,j=1,i>j}^N\} \in \mathbb{R}^{\frac{N \cdot (N-1)}{2} \times D_{\text{sbert}}}, \quad (9.4)$$

where D_{sbert} denotes the embedding size of SBERT extractor. In short, it concatenates all the sentence embedding differences $\mathcal{E}_{\text{sbert}}(Y_i) - \mathcal{E}_{\text{sbert}}(Y_j)$ where $i > j$, resulting in an utterance-level noise embedding $E_{\text{LN}}^{\text{utt}} \in \mathbb{R}^{N \cdot (N-1)/2 \times D_{\text{sbert}}}$. The key idea is, Y_i ranks lower than Y_j in the N-best hypotheses list, which thus presents lower confidence and worse transcription quality, i.e., more *language noise*. Therefore, Eq. (9.4) serves as a measurement of the audio noise in language space. The worse noisy speech would lead to larger ASR decoding uncertainty and thus more diverse N-best hypotheses, so that Eq. (9.4) can capture larger diversity embedding.

Token-Level Noise Embedding Apart from utterance-level embedding, we also propose to extract token-level noise embedding that directly corresponds to the WER metric of ASR task. As shown in the bottom-right part of Fig. 9.2, similar to the calculation of edit distance, we first forced-align the N-best hypotheses to the same length with zero padding (i.e., “ $\emptyset$ ”). The aligned N-best hypotheses $\mathcal{Y}_N^{ali} = \{Y_1^{ali}, Y_2^{ali}, \dots, Y_N^{ali}\}$ clearly illustrates the token difference between different candidatures, where each utterance contains T tokens that comes from ASR vocabulary $\mathcal{V}$ plus zero padding $\emptyset$:

$$Y_i^{ali} = [y_{i_1}^{ali}, y_{i_2}^{ali}, \dots, y_{i_T}^{ali}], \quad y_{i_t}^{ali} \in \mathcal{V} \cup \emptyset, \quad (9.5)$$

Inspired by edit distance, we design an “edit embedding” to capture the token-level difference between two hypotheses, which directly corresponds to their gap in final WER performance. Then, similar to Eq.(9.4), we calculate the token-level noise embedding by summing up the edit embedding between different pairs of hypotheses in the N-best list:

$$\begin{aligned} E_{LN}^{tok} &= \text{Concat}\{E_{\text{edit}}(Y_i^{ali}, Y_j^{ali})_{i,j=1,i>j}^N\} \in \mathbb{R}^{\frac{N(N-1)}{2} \times D_{\text{sbert}}}, \\ E_{\text{edit}}(Y_i^{ali}, Y_j^{ali}) &= \sum_{t=1}^T [\mathcal{E}_{\text{sbert}}(y_{i_t}^{ali}) - \mathcal{E}_{\text{sbert}}(y_{j_t}^{ali})], \end{aligned} \quad (9.6)$$

Note that we employ SBERT again to extract the token embedding, as it can produce informative embeddings for both utterances and tokens [699].

9.3.2 Audio Noise Distillation

After extracting the language-space noise embedding from N-best hypotheses, we further propose an audio noise distillation approach via mutual information estimation to enhance its noise representation ability. Mutual information (MI) is a measure of dependence between random variables based on the Shannon entropy, which is equivalent to the Kullback-Leibler (KL) divergence between the joint distribution and the product of the marginal distribution of random variables. Given two random variables X and Z , their MI can be calculated by:

$$I(X; Z) = D_{KL}(\mathbb{P}_{XZ} \parallel \mathbb{P}_X \mathbb{P}_Z), \quad (9.7)$$

where $D_{KL}(\mathbb{P} \parallel \mathbb{Q})$ denotes KL-divergence. However, it is intractable to directly calculate MI based on Eq. (9.7), so we leverage an estimation method called mutual information neural estimation (MINE) from previous work [56]. MINE employs a

statistics network $\psi_\theta : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}$ parameterized by $\theta \in \Theta$ to estimate a *neural information measure*:

$$I_\Theta(X; Z) = \sup_{\theta \in \Theta} \mathbb{E}_{\mathbb{P}_{XZ}}[\psi_\theta] - \log(\mathbb{E}_{\mathbb{P}_X \mathbb{P}_Z}[e^{\psi_\theta}]), \quad (9.8)$$

In practice, we employ the extracted language-space noise embedding E_{LN} and noisy audio embedding $\mathcal{E}_{\text{ASR}}(X_n)$ as the joint distribution, while using E_{LN} and clean audio embedding $\mathcal{E}_{\text{ASR}}(X_c)$ as the marginal distribution, as the noise information only exists in noisy speech.

9.4 Performance Evaluation

Experiment Setup We conduct experiments on the RobustHP dataset. To verify the general effectiveness of using LLMs for ASR, [330] utilizes various latest LLMs for evaluation, including LLaMA-2-7B/13B [822], LLaMA-7B [819] and Falcon-7B [646]. Here, we present the results based on LLaMA-2-7B. We follow the LLM-Adapter in previous work [989] for both LLM finetuning and noise embedding incorporation.

Metrics We report experimental results in terms of word error rate (WER) and relative GER improvement. We also report two oracle WERs for reference: (1) N-best oracle o_{nb} : WER of the “best candidate” in N-best list, and (2) compositional oracle o_{cp} : best achievable WER using all the tokens in N-best hypotheses. They indicate the upper-bounds of rerank and GER (using occurred tokens), respectively.

Numerical Results Table 9.1 presents the experiment results on LLaMA-2-7B. First, we can observe minor gains of performance brought by typical LM rescoring over the Whisper ASR baseline. Compared to LM rescoring, GER achieves promising progress by leveraging LLMs to generate transcription, while its performance gains in most noisy conditions except CHiME-4 are still limited. Introducing audio denoising further improves the result but suffers from the cross-modality gap. In comparison, with the proposed language-space denoising approach, our RobustGER achieves significant gains of performance in various noise conditions, with up to 53.9% GER improvement in terms of WER metric, where some results even surpass the reranking upper-bound.

Table 9.1 WER (%) results of RobustGER with LLaMA-2-7b finetuning. “LM_{rank}” denotes LM rescore. “+ Audio Denoising” denotes introducing audio embedding to denoise GER. o_{nb} and o_{cp} respectively denote the N-best oracle and compositional oracle. The subscript percentage denotes relative WER reduction over ASR baseline, i.e., GER improvement. Bold value represents best performance among the compared methods

Test Set		Baseline	LM _{rank}	GER	+ Audio Denoising	RobustGER	Oracle	
							o_{nb}	o_{cp}
CHiME-4	<i>test-real</i>	12.6	12.2	6.5 _{-48.4%}	6.4 _{-49.2%}	5.6 _{-55.6%}	10.5	3.0
	<i>test-simu</i>	15.4	14.5	9.2 _{-40.3%}	9.0 _{-41.6%}	8.2 _{-46.8%}	12.9	5.0
	<i>dev-real</i>	10.6	10.3	5.0 _{-52.8%}	4.9 _{-53.8%}	4.1 _{-61.3%}	9.1	2.1
	<i>dev-simu</i>	12.4	11.9	6.8 _{-45.2%}	6.6 _{-46.8%}	5.8 _{-53.2%}	10.6	3.3
	<i>avg.</i>	12.8	12.2	6.9 _{-46.1%}	6.7 _{-47.7%}	5.9 _{-53.9%}	10.8	3.4
VB-DEMAND	<i>baby-cry</i>	8.0	7.8	7.0 _{-12.5%}	6.9 _{-13.8%}	6.0 _{-25.0%}	4.5	3.0
	<i>helicopter</i>	8.4	8.1	7.4 _{-11.9%}	7.3 _{-13.1%}	6.9 _{-17.9%}	4.8	3.2
	<i>crowd-party</i>	22.6	22.3	21.4 _{-5.3%}	21.0 _{-7.1%}	19.2 _{-15.0%}	16.5	11.5
	<i>avg.</i>	13.0	12.7	11.9 _{-8.5%}	11.7 _{-10.0%}	10.7 _{-17.7%}	8.6	5.9
NOIZEUS	<i>babble</i>	16.5	16.7	16.5 _{-0.0%}	16.1 _{-2.4%}	14.5 _{-12.1%}	9.5	5.8
	<i>car</i>	17.4	16.8	15.3 _{-12.1%}	15.2 _{-12.6%}	14.9 _{-14.4%}	9.9	7.9
	<i>station</i>	12.0	11.6	10.3 _{-14.2%}	10.3 _{-14.2%}	9.5 _{-20.8%}	6.6	5.0
	<i>train</i>	15.3	15.2	14.9 _{-2.6%}	15.0 _{-2.0%}	14.9 _{-2.6%}	10.3	7.9
	<i>street</i>	17.4	17.2	17.4 _{-0.0%}	17.1 _{-1.7%}	16.1 _{-7.5%}	12.4	9.9
	<i>airport</i>	11.2	11.0	10.7 _{-4.5%}	10.5 _{-6.3%}	9.5 _{-15.2%}	7.9	4.5
	<i>exhibition</i>	13.2	13.2	12.8 _{-3.0%}	12.4 _{-6.1%}	9.5 _{-28.0%}	8.3	5.8
	<i>restaurant</i>	13.2	13.0	12.4 _{-6.1%}	12.5 _{-5.3%}	12.0 _{-9.1%}	8.7	6.2
	<i>avg.</i>	14.5	14.3	13.8 _{-4.8%}	13.6 _{-6.2%}	12.6 _{-13.1%}	9.2	6.6
LS-FreeSound	<i>metro</i>	9.9	9.8	9.5 _{-4.0%}	9.4 _{-5.1%}	8.9 _{-0.1%}	7.9	4.9
	<i>car</i>	4.0	4.0	3.7 _{-7.5%}	3.5 _{-12.5%}	3.1 _{-22.5%}	3.0	1.8
	<i>traffic</i>	8.3	8.2	8.0 _{-3.6%}	7.8 _{-6.0%}	7.5 _{-9.6%}	6.8	4.5
	<i>cafe</i>	9.8	9.5	8.1 _{-17.3%}	8.1 _{-17.3%}	7.5 _{-23.5%}	7.1	4.6
	<i>babble</i>	32.0	31.8	31.3 _{-2.2%}	31.6 _{-1.3%}	31.1 _{-2.8%}	28.7	19.3
	<i>ac/vacuum</i>	12.4	12.5	12.3 _{-0.8%}	12.1 _{-2.4%}	11.4 _{-8.1%}	10.2	6.2
	<i>avg.</i>	12.7	12.6	12.2 _{-3.9%}	12.1 _{-4.7%}	11.6 _{-8.7%}	10.6	6.9
RATS	<i>test</i>	45.7	45.6	45.2 _{-1.1%}	44.8 _{-2.0%}	43.2 _{-5.5%}	38.8	23.6

Chapter 10

Benchmarking Foundation Models Using Synthetic Datasets



Abstract With the popularity of foundation models, recent years have witnessed a paradigm shift in deep learning from task-centric model design to task-agnostic representation learning and task-specific fine-tuning. Pretrained model representations are commonly evaluated extensively across various real-world tasks and used as a foundation for different downstream tasks. This chapter presents a solution called **SynBench**, as proposed in Ko et al. (What would gauss say about representations? probing pretrained image models using synthetic gaussian benchmarks. In: International Conference on Machine Learning (2024)), for assessing the quality of representations in a task-agnostic way. To circumvent the need for real-world data in evaluation, we explore the use of synthetic binary classification tasks with Gaussian mixtures to probe pretrained vision models and compare the robustness-accuracy performance on pretrained representations with an idealized reference. The approach offers a holistic evaluation, revealing intrinsic model capabilities and reducing the dependency on real-life data for model evaluation.

10.1 Introduction

With the popularity of foundation models, the use of large pretrained neural networks for efficient fine-tuning on downstream tasks has prevailed in many domains such as vision, language, and speech. Instead of designing task-dependent neural network architectures for different downstream tasks, the current methodology focuses on the principle of task-agnostic pretraining and task-specific finetuning. This methodology uses a neural network pretrained on a large-scale broad dataset to extract generic representations of the input data, which we call *pretrained representations* for simplicity. The pretrained representations are then used as a foundation [67] to solve downstream tasks. Prevalent ways include training a linear head (i.e., linear probing) on the representations with the labels provided by a downstream dataset, or simply employing zero-shot inference.

When gauging the usefulness of a pretrained model, it is a convention to conduct evaluations on selected public datasets. For example, ViT [193] reports accuracy on 25 tasks, CLIP [674] on 27 datasets, and PLEX [824] on over 40 datasets

to systematically evaluate different reliability dimensions on both vision and language domains. However, this convention has several drawbacks. For example, the evaluation process evidently poses significant computational overhead on the model trainer and raises data privacy concerns, setting a high bar for new model designs and large-scale AI governance. More importantly, the evaluation result is dependent on specific evaluation datasets. Thus the nominal evaluation score can be inconclusive if the evaluation data are biased or under-representative. For instance, ViT-L/16 is reportedly performing better than ViT-B/16 on 23 out of 27 linear probing tasks according to [674, Table 10], but worse than ViT-B/16 on FoodSeg103 [913, Table 8], X-ray images [616, Table 4-8], and magnetic resonance imaging [835, Table 2-3] tasks. Fundamentally, a poor probing result might come from either (1) evaluation data bias, (2) true model deficiency, or both. In [414], the authors attempt to disentangle the effect of the two and focus on designing well-posed sanity checks for the latter. They utilize synthetic data generated from class-conditional data prior, whose optimal classification strategy is known, and compare the optimal strategy with representations' linear separability. For example, Fisher's linear discriminant rule [380, 651] decides the optimal strategy for Gaussian distribution. If the data can be separated with 90% accuracy in the raw input space and 60% in the representation space, then the pretrained model has an intrinsic deficiency. Building on that, the trending practice of pretraining and fine-tuning also signifies immediate damage to all adapted applications if the foundation model has hidden risks [67], such as lacking robustness to adversarial examples. These types of risks may not be informed by the standard accuracy as they do not correlate well [785]. Luckily, similar to Fisher's linear discriminant rule for the optimal standard accuracy, [172] has characterized the optimal classification strategy in the presence of input perturbations. The sanity check can thereby evaluate the robustness of pretrained models by considering the same synthetic conditional Gaussian data prior.

Besides being great candidates for establishing well-posed problems, the idea of probing foundation models with synthetic conditional Gaussians is also motivated by the longstanding practice of Gaussian modeling in signal processing [287], data mining [291], machine learning [407, 837, 1051], and other engineering fields. For example, Gaussian mixtures have found applications in modeling noise, magnetic field inhomogeneities, biological variations of tissues in magnetic resonance imaging [686], and computerized tomography [727]. The facts that Gaussian mixture models often lead to mathematically tractable problems [547, 585, 698] and the abundance of analytical tools available for Gaussian models [1, 172, 380, 658] also inspire the study on how Gaussian mixtures can be leveraged for evaluating pretrained image models.

An ideal pretrained model should entail both good accuracy and robustness, and the level of goodness is desired to be measurable in a task/data-agnostic manner. In this chapter, we present *SynBench* [414] to precisely address this requirement. Specifically, *SynBench* establishes a theoretical reference characterizing the robustness-accuracy trade-off of the synthetic data based on the Bayes optimal linear classifiers. Then, *SynBench* obtains the representations of the same synthetic data




Fig. 10.1 Overview of SynBench [414]. **Step 1:** generate class-conditional Gaussian and form the inputs to the pretrained model; **Step 2:** gather rendered representations; **Step 3:** measure the expected robustness bound under a range of threshold accuracy for both input synthetic data and their representations according to Eq. (10.2) and obtain the expected bound-threshold accuracy plot; **Step 4:** calculate SynBench score by the relative area under the curve of the representations (area B) to the inputs (area A + area B) in the expected bound-threshold accuracy plot. The closer the ratio is to 1, the better the quality of pretrained representations is, in terms of the robustness-accuracy characterization

from the pretrained model and compares them to the reference. Finally, we define the ratio of area-under-the-curves in robustness-accuracy plots, *SynBench-Score*, as a quantifiable metric of the pretrained representation quality. The entire procedure of SynBench is illustrated in Fig. 10.1.

We discuss possible use case of SynBench as follows. We view SynBench as a necessary and minimum model test in the sense that, with perfect data sampled from an ideal distribution, any undesirable deteriorated behavior (such as weakened robustness) reveals the weaknesses of the representation model that could possibly lead to vulnerabilities in real-life downstream tasks. Therefore, in designing this minimum test, it is important that the task has a theoretical ideal (and optimal) solution (i.e. the trade-off preserved by class conditional Gaussians as derived in [414]). Here are some possible scenarios to use SynBench:

- **Model auditing:** use SynBench to generate diverse psuedo tasks (e.g., with different difficulty levels) and compare them with theoretically optimal results, for a comprehensive evaluation on the capability of a pre-trained model
- **Hyperparameter tuning:** as shown in [414], SynBench can be used for hyperparameter selection in robust linear probing, which leads to improved performance in the considered downstream tasks.
- **Model selection (without using downstream data):** without the knowledge of downstream applications, one can use SynBench to rank the quality of pre-trained representations [414]. It is also possible to incorporate some known statistics of the downstream dataset into guided synthetic data generation and evaluation in SynBench, as discussed in [414].
- **Model training:** while updating a model in the pre-training state, one can use SynBench to ensure the model performance (in terms of SynBench-Score) is aligned.

In addition to using Gaussian synthetic data to benchmark vision foundation models in image classification, the authors also extend the SynBench framework to characterize the robustness-accuracy of language models in [413].

10.2 Background and Related Work

Pretrained Models in Vision In the past few years, much focus in the machine learning community has been shifted to training representation networks capable of extracting features for a variety of downstream tasks with minimal fine-tuning. Nowadays, many common vision tasks are achieved with the assistance of good backbones, e.g. classifications [119, 193, 228, 905, 919, 963], object detection [527, 696], segmentation [117, 917], etc. Among the popular backbones, vision transformers (ViT) [193] and convolutional models (e.g. ResNet [290]) have attracted enormous interest. We will exemplify the use of SynBench using several pretrained ViTs and ResNets.

Benchmarking Pretrained Models Since pretrained models are used as a foundation for different downstream tasks, it is central to transfer learning [608, 663], and also tightly related to model generalization [93, 667]. To benchmark the performance of a pretrained model, it is a convention to apply the pretrained model for a number of popular tasks and conduct linear probing on the representations [119, 129, 133, 193]. Besides accuracy-based probing methods, evaluation methods have been proposed based on information theory and minimum description length [66, 848], surplus description length [899], maximum evidence [959], Fisher discriminant analysis [747], among others. These metrics are reliant on the label information of the downstream tasks and are hence task-specific.

Lately, more fundamental questions related to pretrained models are brought up [67, 753, 824, 986]. Bommasani et al. [67] raised practical concerns about the homogenization incentivized by the scale of the pretraining. Although homogenization might help in achieving competitive performance for some downstream tasks, the defects are also inherited by all these downstreams. On that account, a more careful study of the fundamentals of pretrained models is of paramount importance. Tran et al. [824] explored the reliability of pretrained models by devising 10 types of tasks on 40 datasets. It is further pointed out by Zhang and Ré [986] in 9 benchmarks that pretrained models may not be robust to subpopulation or group shift. The adversarial robustness is benchmarked by Shao et al. [746] and Paul and Chen [644].

Optimal Representations In the seminal work of deep representation theory, [10] depicted the desired optimal representations in supervised learning to be sufficient for the downstream task, invariant to the effect of nuisances, maximally disentangled, and have minimal mutual information between representations and inputs. Focusing more on generalization than compression, [198] provided the optimal representation based on $\mathcal{V}$ -information [930]. Ruan et al. [716] defined the optimal

representations for domain generalization. Dubois et al. [197] characterized idealized representations in self-supervised learning as ones that are well-distinguished by the desired family of probes for potential invariant tasks, have sufficiently large dimensions, and be invariant to input augmentations.

Why SynBench? To enable quantifying representation quality in the pretraining stage, SynBench differs from the above frameworks as it does not need knowledge of any real-world downstream data. Moreover, SynBench has full control of the evaluation set via synthetic data generation. With the assumed synthetic data distribution, we can theoretically characterize the reference robustness-accuracy trade-off. Therefore, SynBench provides a standardized quality metric with theoretical groundings and evaluates for representations induced by pretrained models at a low cost.

10.3 SynBench

Without the knowledge of the downstream tasks and data, SynBench [414] aims to develop a task-agnostic framework to evaluate some fundamental behaviors of the representation network. In this chapter, we inspect and quantify how representation networks preserve the robustness and accuracy enjoyed by the original synthesized data. On the whole, we measure the idealized robustness-accuracy trade-off using synthetic data. By propagating the Gaussian realizations through different representation networks, we can also compare the robustness-accuracy trade-off for representations. We start this section by giving the preliminaries on the synthetic data of interest.

10.3.1 Synthetic Data

We consider binary classification problems with data pair (x, y) generated from the mixture of two Gaussian distributions $P_{\mu_1, \mu_2, \Sigma}$, such that $x|y = 1 \sim \mathcal{N}(\mu_1, \Sigma)$, $x|y = -1 \sim \mathcal{N}(\mu_2, \Sigma)$, or equivalently,

$$x - \frac{\mu_1 + \mu_2}{2} | y \sim \mathcal{N}(y\tilde{\mu}, \Sigma), \quad (10.1)$$

where $y \in \mathcal{C} = \{+1, -1\}$, $P(y = +1) = \tau$, $P(y = -1) = 1 - \tau$, and $\tilde{\mu} = \frac{\mu_1 - \mu_2}{2}$. We focus on the class-balanced case ($\tau = \frac{1}{2}$). The imbalanced case is discussed in [414]. When sampling from this idealized distribution, we eliminate the factor of data bias and can test the accuracy and robustness degradation in an ideal setting.

Let $\|\cdot\|_p$ denote the ℓ_p norm of a vector for any $p \geq 1$. For a given classifier f and input x with $f(x) = y$, where y is the predicted label, it is not

rational for the classifier to respond differently to $x + \delta$ than to x for a small perturbation level measured by $\|\delta\|_p$, i.e. inconsistent top-1 prediction [261, 796]. Therefore, the level of (adversarial) robustness for a classifier can be measured by the minimum magnitude of perturbation that causes misclassification, i.e. $\|\Delta\|_p := \min_{\delta: f(x+\delta) \neq f(x)} \|\delta\|_p$. For a generic function f , solving the optimization problem exactly is hard [389, 763]. Luckily, one can readily solve for the optimization if f is affine [597].

10.3.2 Main Theorem

In what follows, we will leverage this point and focus on the linear classifier that minimizes robust classification error. An ideal candidate classifier for the class conditional Gaussian in (10.1) is specified by the robust Bayes optimal classifier [59, 188]. Specifically, it is stated that the optimal robust classifier (with a robust margin ϵ) for data generated from (10.1) is a linear classifier. We derive the following result as a direct application of the fact. To simplify the exposition, we focus on the ℓ_2 norm in the remainder of this paper. The general ℓ_p -norm results are given in [414]. We use “bound” to denote the minimal perturbation of a sample. We first formally state the theorem in [414] that serves as the foundation of our SynBench framework.

Theorem 10.1 ([414]) *For any sample x , the optimal robust classifier f_ϵ for $P_{\mu_1, \mu_2, \Sigma}$ gives*

- (i) *the bound (decision margin)*

$$\|\Delta\|_2 = \frac{|(x - \frac{\mu_1 + \mu_2}{2})^T \Sigma^{-1} (\tilde{\mu} - z_\Sigma(\tilde{\mu}))|}{\|\Sigma^{-1} (\tilde{\mu} - z_\Sigma(\tilde{\mu}))\|_2},$$
- (ii) *the scaled bound* $\|\tilde{\Delta}\|_2 = \frac{|(x - \frac{\mu_1 + \mu_2}{2})^T \Sigma^{-1} (\tilde{\mu} - z_\Sigma(\tilde{\mu}))|}{|\tilde{\mu}^T \Sigma^{-1} (\tilde{\mu} - z_\Sigma(\tilde{\mu}))|}.$

For a sample $x \sim P_{\mu_1, \mu_2, \Sigma}$, it further gives

- (iii) *the standard accuracy* $a = \Phi(\frac{\tilde{\mu}^T \Sigma^{-1} (\tilde{\mu} - z_\Sigma(\tilde{\mu}))}{\|\Sigma^{-1} (\tilde{\mu} - z_\Sigma(\tilde{\mu}))\|_\Sigma}),$
- (iv) *the expected scaled bound of correct samples*

$$\mathbb{E}[\|\tilde{\Delta}\|_2 \mid f_\epsilon(x) = y] = \frac{1}{\sqrt{2\pi}} \frac{1}{a \Phi^{-1}(a)} e^{-\frac{1}{2}(\Phi^{-1}(a))^2} + 1,$$

where z_Σ is the solution of the convex problem $\arg \min_{\|z\|_2 \leq \epsilon} (\tilde{\mu} - z)^T \Sigma^{-1} (\tilde{\mu} - z)$ and Φ denotes the CDF of the standard normal distribution.

We note that for samples drawn from $P_{\mu_1, \mu_2, \Sigma}$, $\Sigma = \sigma^2 I_d$, all ϵ -robust Bayes optimal classifier overlap with each other. For a general covariance Σ , the ϵ of an ϵ -robust Bayes classifier specifies the desired size of margin and demonstrates the robustness accuracy trade-off. We give an illustrative 2D class-conditional Gaussian example in Fig. 10.2a, where different ϵ -robust Bayes classifiers give different overall margins at the cost of accuracy. As ϵ increases, the robust Bayes optimal




Fig. 10.2 Illustration of robustness-accuracy trade-off suggested by ϵ -robust Bayes optimal classifiers. (a) depicts a class-conditional 2D Gaussian case with decision boundaries drawn by ϵ -robust Bayes optimal classifiers of varying ϵ values. (b) draws the theoretically characterized robustness-accuracy trade-off given in Theorem 10.1(iv)

classifier rotates counterclockwise, leading to increased misclassifications, but also overall enlarged margins.

10.3.3 Objective

For a given representation network parameterized by θ , we are interested in evaluating the expected bounds on synthetic data and their representations, under a thresholding accuracy a_t . That is, $\mathbb{E}_{\mu \sim \mathbb{P}_\mu, \Sigma \sim \mathbb{P}_\Sigma, x - \bar{\mu} | y \sim \mathcal{N}(y\mu, \Sigma)} [\|\bar{\Delta}\|_2 \mid f_\epsilon(x) = y, a > a_t]$ for $\bar{\Delta} = \bar{\Delta}_x$ and $\bar{\Delta}_z$, where $\mathbb{P}_\mu$ and $\mathbb{P}_\Sigma$ characterize the probability density function of the synthetic data manifold of interest, $\bar{\mu}$ is a translation vector allowing non-symmetric class conditional Gaussian, and $\bar{\Delta}_x$ and $\bar{\Delta}_z$ denote the bounds on synthetic data and representations respectively. Here, without the prior of applications, we assume $\mu = s \cdot 1_d / \sqrt{d}$, where s denotes a random variable that follows uniform distribution and $1_d / \sqrt{d}$ is the normalized all-ones vector. For simplicity, we let $\Sigma = I_d$. Formally, we define the accuracy-constrained expected bound $E_{\theta, \epsilon}(a_t)$ as

$$\begin{aligned} E_{\theta, \epsilon}(a_t) &= \mathbb{E}_{s, x} [\|\bar{\Delta}\|_2 \mid f_\epsilon(x) = y, a(s, \epsilon) > a_t] \\ &= \sum_i \mathbb{E}_x [\|\bar{\Delta}\|_2 \mid f_\epsilon(x) = y] \mathbb{1}_{a(s_i, \epsilon) > a_t} p(s_i), \end{aligned} \quad (10.2)$$

where $\mathbb{1}_{a(s_i, \epsilon) > a_t}$ is the indicator function specifying the s_i, ϵ -dependent accuracy a that surpasses the threshold accuracy a_t . The detailed derivation can be found in [414]. In the following sections, we will illustrate how to calculate the inner expectation term $\mathbb{E}_x [\|\bar{\Delta}\|_2 \mid f_\epsilon(x) = y]$ for both the raw data (synthetic data) and representations.

Raw Data For raw data synthesized from $P_{\mu_1, \mu_2, \Sigma}$ according to (10.1), the inner expectation term is given by Theorem 10.1(iv) $\mathbb{E} [\|\bar{\Delta}_x\|_2 \mid f_\epsilon(x) = y] = \frac{1}{\sqrt{2\pi}} \frac{1}{a\Phi^{-1}(a)} e^{-\frac{1}{2}(\Phi^{-1}(a))^2} + 1$, where a denotes the standard accuracy. The subscript x in the expected scaled bound $\mathbb{E} [\|\bar{\Delta}_x\|_2 \mid f_\epsilon(x) = y]$ indicates the raw data space, to distinguish from the scaled bound to be derived for representations. We highlight that Theorem 10.1(iv) directly shows a robustness-accuracy trade-off. We plot the expected scaled bound as a function of accuracy in Fig. 10.2b, which holds true when the data follow (10.1) exactly. In SynBench, we treat this theoretically-derived robustness-accuracy trade-off as the reference, enabling a fair comparison among representations induced by different pretrained models.

Representations Given a pretrained network, we gather the representations of the Gaussian realizations and quantify the bound induced by robust Bayes optimal classifier in the representation space. When deriving the robust Bayes optimal classifier, we model the representations by a general conditional Gaussian $z|y =$

$1 \sim \mathcal{N}(\mu_1, \Sigma)$, $z|y = -1 \sim \mathcal{N}(\mu_2, \Sigma)$. By Theorem 10.1(ii), we consider the optimal robust classifier for the modeled conditional Gaussian in the representation space to calculate the scaled bound $\|\bar{\Delta}_z\|_2 = \frac{|(z - \frac{\mu_1 + \mu_2}{2})^T \Sigma^{-1}(\bar{\mu} - z_\Sigma(\bar{\mu}))|}{|\bar{\mu}^T \Sigma^{-1}(\bar{\mu} - z_\Sigma(\bar{\mu}))|}$ for correctly-classified samples and the inner expectation is estimated empirically. It should be noted that now the Bayes optimal classifier does not necessarily coincide with the robust Bayes optimal classifier even when we synthesized the dataset with an identity matrix covariance in the input space.

10.3.4 Robustness-Accuracy Quantification

Recall that we aim to calculate $E_{\theta, \epsilon}(a_t) = \sum_i \mathbb{E}_{x|y \sim \mathcal{N}(y s_i \cdot 1_d / \sqrt{d}, I_d)} [\|\bar{\Delta}\|_2 | f_\epsilon(x) = y] \cdot \mathbb{1}_{a(s_i, \epsilon) > a_t} p(s_i)$ for both raw data and the representations (i.e. $\|\bar{\Delta}_x\|$ and $\|\bar{\Delta}_z\|$). We treat the expected bounds of the raw data under a threshold accuracy as the reference. Given a representation network, we compare the expected bounds of the representations rendered by representation networks with the reference.

In the implementation of Synbench [414], the authors take $s \sim \mathcal{U}\{0.1, 5\}$ under the guidance of Theorem 10.1(iii). Specifically, as Theorem 10.1(iii) gives an analytical expected accuracy for class conditional Gaussian, we can obtain the desired range of s by giving the accuracy. Since we are interested in having the reference as a class conditional Gaussian that yields accuracy from 55% to almost 100%, we set the starting and ending s by the fact that $\Phi(0.1) \approx 0.55$ and $\Phi(5) \approx 1.0$. As discussed in [414], with more accurate modeling of the data manifold of interest, SynBench can give a more precise capture of the pretrained representation performance.

When the data is perfect Gaussian (e.g. input synthetic data), we calculate $E_{\theta, \epsilon}(a_t)$ as detailed in Sect. 10.3.3. We note that $\bar{\Delta}_x$ is independent of pretrained network parameters θ , and all the ϵ -robust classifiers f_ϵ in the input space overlap with each other when $\Sigma = I_d$. We hereby denote the desired metric on the input synthetic data by $E(a_t)$, to distinguish from that on the representations $E_{\theta, \epsilon}(a_t)$. For representations, we calculate $E_{\theta, \epsilon}(a_t)$ following Sect. 10.3.3 and the expectation is estimated empirically. We show an example of the probing results in Fig. 10.3.

To integrate over all the desired threshold accuracy, we use the area under the curve (AUC) and give the ratio to the reference by

$$\text{SynBench-Score}(\theta, \epsilon, a_T) = \frac{\int_{a_T}^1 E_{\theta, \epsilon}(a_t) da_t}{\int_{a_T}^1 E(a_t) da_t}, \quad (10.3)$$

which correspond to the relative area $\frac{\text{area B}}{\text{area A} + \text{area B}}$ in Fig. 10.3. Values of SynBench-Score closer to 1 imply better probing performance on pretrained representations. To summarize, SynBench framework generates a sequence of proxy tasks with different difficulty levels (monitored by s). With each proxy task, we can obtain an accuracy




Fig. 10.3 An example of the robustness-accuracy quantification of representations for ViT-B/16. (Left) The expected bound-threshold accuracy plot for the input raw data ($E(a_t)$) and representations ($E_{\theta, \epsilon}(a_t)$) with $\epsilon = 0 \sim 0.8$. (Right) To calculate the SynBench-Score for $\epsilon = 0$ (top) and $\epsilon = 0.6$ (bottom), we use the definition $\text{SynBench-Score}(\theta, \epsilon, a_t) = \frac{\text{area B}}{\text{area A} + \text{area B}}$ (refer to (10.3)), which gives $\text{SynBench-Score}(\theta_{\text{ViT-B/16}}, 0, 0.7) = 0.33$ and $\text{SynBench-Score}(\theta_{\text{ViT-B/16}}, 0.6, 0.7) = 0.20$

and an expected bound (Sect. 10.3.3). With gathered pairs of accuracy and expected bound, we filter ones whose accuracy is below a threshold accuracy (x-axis), and calculate the accuracy-constrained expected bound to reflect the robustness level (y-axis). With this, the AUC will counter the discriminative power of the foundation model given an idealized distribution, as well as the robustness level.

10.4 Performance Evaluation

Experiment Setup We will calculate SynBench-Scores for pretrained models and make pair-wise comparisons. For example, ViT-B/16 is a fine-tuned pretrained model from ViT-B/16-in21k. By checking their SynBench-Scores, we could understand how the fine-tuning procedure helps or worsens the performance. In order to systematically understand how each network attribute affects the robustness-accuracy performance, it is desirable to control the variates. We list and compare 10 pretrained vision transformers (ViTs) [94, 133, 193] and ResNets [130] in Table 10.1.

Baselines We refer to recent work [747, 899, 959] and report the validation accuracy (Val loss), minimum description length (MDL), surplus description length (SDL), logarithm of maximum evidence (LogME) and self-challenging Fisher discriminant analysis (SFDA), following the official implementation from the literature on our synthetic proxy task as baselines [747, 899].

Table 10.1 Model descriptions. The performance of models might be nuanced by scheduler, curriculum, and training episodes, which are not captured in the table

Model	Arch.	Pretraining	Fine-tuning	Patch	# parameters (M)
ViT-Ti/16	ViT-Tiny	Imgn21k	Imgn1k	16	5.7
ViT-B/16	ViT-Base	Imgn21k	Imgn1k	16	86.6
ViT-B/16-in21k	ViT-Base	Imgn21k	No	16	86.6
ViT-L/16	ViT-Large	Imgn21k	Imgn1k	16	304.3
ViT-S/16-DINO	ViT-Small	self-Imgn1k	No	16	21.7
ViT-S/8-DINO	ViT-Small	self-Imgn1k	No	8	21.7
ViT-B/16-DINO	ViT-Base	self-Imgn1k	No	16	85.8
ViT-B/8-DINO	ViT-Base	self-Imgn1k	No	8	85.8
Resnet50-SimCLRv2	Resnet50	self-Imgn1k	No	–	144.4
Resnet101-SimCLRv2	Resnet101	self-Imgn1k	No	–	261.2
<i>Variation</i>					
Model size	ViT-{Ti,B,L}/16, ViT-{S,B}/16-DINO, ViT-{S,B}/8-DINO, Resnet{50,101}-SimCLRv2				
Finetuning	ViT-B/16{,-in21k}				
ViT patch size	ViT-S/{16,8}-DINO, ViT-B/{16,8}-DINO				

Evaluation In essence, we expect these real-data-free evaluations for pretrained models can give meaningful performance assessments of possible downstream tasks. For this purpose, we take an average of the accuracy in 27 downstream tasks (cf. [674], Table 10) as in the literature [193, 216, 474, 674, 963] to give a sense of the general performance on possible downstream tasks, and report the Pearson correlation coefficients with SynBench-Scores. Building on top of these, we also show the consistency of SynBench suggestions given different numbers of synthetic realizations compared to the baselines. To provide a comprehensive evaluation, we give $\text{SynBench-Score}(\theta, \epsilon, a_t)$ with a_t ranging from 0.7 to 0.9, and ϵ from 0 to 0.8. Besides the SynBench-Score, we will also report the standard accuracy (SA) and robust accuracy against adversarial perturbations (RA) for studying robustness-accuracy performance.

Numerical Results We list the SynBench-Score of the 10 pretrained representations with their standard and robust accuracy on the class-conditional Gaussian proxy task in Table 10.2. The robust accuracy is obtained by ℓ_2 PGD attack [563] with attack strength 0.2. By referring to rows “ViT-B/16” and “ViT-B/16-in21k”, we see that SynBench will suggest ViT-B/16 over ViT-B/16-in21k, implying that the fine-tuning is beneficial on ViT-B/16-in21k—both networks are pretrained on Imagenet 21k with supervision, whereas ViT-B/16 is further finetuned on Imagenet 1k. We can also use SynBench to evaluate the effect of model sizes. Specifically, we refer to rows “ViT-Ti/16”, “ViT-B/16”, “ViT-L/16”, and see that ViT-B/16 and ViT-L/16 score much higher than ViT-Ti/16, suggesting larger models have

Table 10.2 The SynBench-Score of pretrained representations and the standard/robust accuracy (SA/RA) (%) of their linear probing classifier on class-conditional Gaussian data

Models	SynBench-score ($\epsilon = 0$)	SA	RA
ViT-Ti/16	0.01	76.0	50.8
ViT-B/16	0.33	96.4	52.9
ViT-B/16-in21k	0.20	92.1	51.3
ViT-L/16	0.26	96.1	52.9
ViT-S/16-DINO	0.48	97.9	55.5
ViT-B/16-DINO	0.55	99.3	50.4
ViT-S/8-DINO	0.40	95.8	51.1
ViT-B/8-DINO	0.50	98.8	49.6
Res50-SimCLRv2	0.66	99.8	50.1
Res101-SimCLRv2	0.60	99.4	51.6

better capacities for robustness and accuracy. It is noticeable that ViT-B/16 is generally on par with ViT-L/16 when we vary ϵ [414]. Similar conclusions can also be drawn by referring to self-supervised pretrained representations, rows “ViT-S/-DINO” and “ViT-B/-DINO”. Moreover, if we check rows “ViT-B/16” and “ViT-B/16-DINO”, we compare two pretrained models of the same architecture but trained under different regimes, either supervised or self-supervised. Between these two models, SynBench favors self-supervised trained “ViT-B/16-DINO”, echoing with the inductive bias of self-supervised contrastive learning discovered in recent literature [284].

We run baseline evaluations for the synthetic classification task on pretrained models with dataset size n being 2048, 8192, 32768. Throughout our experiments, we use 2048 test samples in the synthetic dataset. For Val loss, MDL, and SDL, ϵ SC, the smaller the better; for LogME, SFDA, SynBench, the bigger the better. In Fig. 10.4, we illustrate how the correlation between task-agnostic evaluation metrics and real-life data tasks varies with the dataset size n . Specifically, we calculate the Pearson correlation coefficients between the average accuracy in downstream tasks to scores given by Val loss, MDL, SDL, ϵ SC, LogME, SFDA, and SynBench (SDL and ϵ SC are excluded from the figure since they fail to give concrete numbers for small dataset sizes). With 2k synthetic samples, SynBench gives 0.79, whereas Val loss, MDL, LogME, and SFDA range between 0.46 and 0.55; with 8k synthetic samples, SynBench gives 0.89, whereas Val loss, MDL, LogME, and SFDA range between 0.65 and 0.81, surpassing the correlation by vanilla out-of-distribution accuracy (ImageNet-c’s 0.64 and ImageNet-a’s 0.57); with over 30k synthetic samples, Val loss, MDL, and SynBench all indicate very strong correlation (> 0.9) with real-life data accuracy, confirming the feasibility of probing pretrained representations in a task-agnostic yet effective way.




Fig. 10.4 Pearson correlation between task-agnostic metrics (Val loss, MDL, SynBench, LogME, SFDA) and task-specific metrics (the average accuracy on 27 real-life tasks) as functions of the dataset size. Two dashed lines characterize the correlation by transfer datasets' accuracy

Chapter 11

Machine Unlearning for Foundation Models



Abstract This chapter aims to provide a comprehensive understanding of emerging machine unlearning (MU) techniques in foundation models. These techniques are designed to precisely evaluate the impact of specific data and high-level knowledge concepts on model performance and to efficiently and effectively eliminate their (possibly harmful) influence within a pre-trained model in response to users' removal requests. Initially proposed to address data privacy concerns in compliance with the 'right to be forgotten' regulation, MU has become increasingly crucial with the advent of foundation models, as re-training from scratch (after removing the undesired training points) is prohibitively expensive in terms of time, compute, and money. In this section, we explore MU from several key perspectives: foundational concepts and formulations, optimization techniques, adversarial evaluation methods, and practical applications.

Keywords Foundation model · Machine unlearning · Diffusion model

11.1 Introduction

MU (machine unlearning) was initially proposed to address data privacy concerns, particularly in compliance with the “right to be forgotten” regulation [310]. This concept involves reversing the learning process to remove the influence of specific data points in a machine learning (ML) model, thereby preventing information leakage about private data after training is completed [72, 89, 609]. The most straightforward and optimal unlearning approach is known as “exact unlearning”, which involves retraining the ML model from scratch using the remaining training set after removing the data points to be scrubbed.

However, given the costly and prolonged training periods of foundation models, retraining these models to eliminate undesirable data effects is often impractical [526]. As a result, the development of *approximate but faster unlearning* methods has become a significant focus in research. At the same time, the application landscape of MU has rapidly evolved, demonstrating its potential in various areas. For instance, the tension between data owners (e.g., authors) and large language

models (LLMs) service providers is escalating, leading to legislation and legal disputes involving companies like OpenAI, Meta, and the New York Times. This trend is likely to persist due to increasing societal concerns about AI data usage. To address this problem, LLM unlearning is uniquely suited for the removal of copyright-protected content. Unlearning can also be employed to eliminate harmful behaviors, such as the production of toxic, discriminatory, illegal, or morally undesirable outputs e.g., instructions for building CBRN (chemical, biological, radiological, and nuclear) weapons [478].

The MU techniques can be broadly categorized into two types based on the type of ML model they focus on: *MU for discriminative models* (e.g., image classifiers) [214, 259, 354, 369, 809, 881] and *MU for generative models* (e.g., diffusion models for text-to-image generation and LLMs) [205, 214, 241, 360, 370, 427, 478, 566, 948, 960, 978, 981, 1012]. In the following sections, we will detail the objectives and setup of MU and introduce a generic formulation of the MU problem. Subsequently, we will explore MU techniques for both discriminative models and generative models in detail.

11.2 Research Objective, Formulation, and Related Work

The objective of MU is to negate the impact of a specific subset of training data points or a higher-level knowledge concept on a pre-trained model, while preserving its utility for data not subject to unlearning. For a concrete setup of MU, consider the training dataset $\mathcal{D} = \{\mathbf{z}_i\}_{i=1}^N$, consisting of N data samples. Each sample $\mathbf{z}_i$ includes a feature vector $\mathbf{x}_i$ and a possible label $\mathbf{y}_i$ for supervised learning. Let $\mathcal{D}_f \subseteq \mathcal{D}$ represent the subset of data targeted for unlearning, with its complement, $\mathcal{D}_r = \mathcal{D} \setminus \mathcal{D}_f$, being the dataset to retain. We refer to $\mathcal{D}_f$ as the *forget set* and $\mathcal{D}_r$ as the *retain set*, respectively. Prior to unlearning, we have access to an initial model, denoted by θ_o , which has been trained on the full dataset $\mathcal{D}$ using methods like empirical risk minimization (ERM).

Given the above setup, Retrain, an exact yet expensive unlearning approach, entails retraining the model θ_o from scratch, exclusively utilizing the retain set $\mathcal{D}_r$. It is typically regarded as the gold standard in MU [369, 809]. However, due to the prolonged training time and the high cost, Retrain is often impractical. Consequently, approximate unlearning methods have emerged as efficient alternatives. Their objective is to efficiently create an *unlearned model*, denoted as θ_u , leveraging prior knowledge of θ_o and the forget set $\mathcal{D}_f$ and/or the retain set $\mathcal{D}_r$. Following the conceptual framework of MU in [526], the optimization problem to obtain θ_u can be expressed as

$$\theta_u = \arg \min_{\theta} \ell_{\text{MU}}(\theta) := \ell_r(\theta; \mathcal{D}_r) + \lambda \ell_f(\theta; \mathcal{D}_f), \quad (11.1)$$

Table 11.1 Overview of MU methods highlighting differences in relabeling-based forget loss, necessity of random re-initialization, partial model updates during unlearning, and the retain-forget regularization parameter λ within (11.1)

Method	Relabeling	Random re-initialization	Partial model update	$\lambda = 0$
Retrain	✗	✓	✗	✓
FT [881]	✗	✗	✗	✓
EU- k [257]	✗	✓	✓	✓
CF- k [257]	✗	✗	✓	✓
SCRUB [430]	✗	✗	✗	✗
ℓ_1 -sparse [369]	✗	✗	✗	✗
RL [259]	✓	✗	✗	✗
BE [118]	✓	✓	✗	✗
BS [118]	✓	✗	✗	✗
SalUn [214]	✓	✗	✓	✗

where ℓ_f and ℓ_r represent the forget loss and the retain loss, respectively, with $\lambda \geq 0$ acting as a regularization parameter. For instance, fine-tuning using the retain set $\mathcal{D}_r$ equates to setting $\lambda = 0$, aimed to impose catastrophic forgetting of over $\mathcal{D}_f$ after model fine-tuning. The objective functions ℓ_r and ℓ_f can be specified as the training loss and its negative variant, respectively. It is also important to note that the specifics of the problem (11.1) can vary depending on the application scenario, as detailed below.

MU for Image Classification This is the most commonly studied MU problem in the literature [744]. Depending on the composition of the forgetting dataset $\mathcal{D}_f$, MU for image classification can be further categorized into two scenarios: *class-wise forgetting* and *random data forgetting*. The former aims to eliminate the influence of training data points from the same image class, while the latter aims to remove the influence of randomly selected data points from the entire training set.

Evaluating the effectiveness of MU for image classification has involved the use of various metrics. While a consensus is still lacking, we adhere to the recent approach proposed by Jia et al. [369], which considers a comprehensive ‘full-stack’ MU evaluation. This includes *unlearning accuracy (UA)*, given by $1 - \text{accuracy}$ of an unlearned model θ_u on $\mathcal{D}_f$, *membership inference attack (MIA)* on $\mathcal{D}_f$ (to determine whether a sample in $\mathcal{D}_f$ is correctly predicted as a non-training data), *remaining accuracy (RA)*, i.e., the fidelity of an unlearned model θ_u on the remaining training set $\mathcal{D}_r$, *testing accuracy (TA)*, i.e., the generalization of θ_u , and *run-time efficiency (RTE)*, i.e., the computation time of applying an MU method.

Assisted by (11.1), we provide an overview of 9 existing (approximate) unlearning methods examined in this study; see Table 11.1 for a summary. These methods can be roughly categorized into two main groups based on the choice of the forget loss ℓ_f : *relabeling-free* and *relabeling-based*. The latter, relabeling-based methods, assign an *altered* label, distinct from the true label, to the data point targeted for forgetting. Consequently, minimizing ℓ_f compels the unlearned model to discard the

accurate label of the points to be forgotten. These methods include random labeling (**RL**) [259], boundary expanding (**BE**) [118], boundary shrinking (**BS**) [118], and saliency unlearning (**SalUn**) [214]. In contrast, relabeling-free methods utilize fine-tuning on the retain set $\mathcal{D}_r$ to induce catastrophic forgetting or apply gradient ascent on the forget set $\mathcal{D}_f$ to achieve the forgetting objective. These methods include fine-tuning (**FT**) [881], exact unlearning restricted to the last k layers (**EU- k**) [257], catastrophically forgetting the last k layers (**CF- k**) [257], scalable remembering and unlearning unbound (**SCRUB**) [430] and ℓ_1 -**sparse** MU [369].

MU for Image Generation by Diffusion Models (DMs) Recent studies have demonstrated that well-trained DMs, such as stable diffusion (SD) [709], can generate images containing harmful content, such as ‘nudity’, when subjected to inappropriate text prompts [732]. This has raised concerns regarding the safety of DMs. To this end, current solutions endeavor to compel DMs to effectively *erase* the influence of inappropriate text prompts in the diffusion process, e.g., referred to as *concept erasing* in [241] and *learning to forget* in [978]. These methods are designed to thwart the generation of harmful image content, even in the presence of inappropriate prompts. The pursuit of safety improvements for DMs aligns with the concept of MU [89, 609, 744, 809, 925]. The MU’s objective of achieving ‘the right to be forgotten’ makes the current safety enhancement solutions for DMs akin to MU designs tailored for the specific context of DMs. In light of this, we refer to DMs developed with the purpose of eliminating the influence of harmful prompts as *unlearned DMs*.

For ease of understanding, we briefly review the diffusion process and DM training. Let $\epsilon_\theta(\mathbf{x}_t|c)$ symbolize the noise generator parameterized by θ , conditioned on the text prompt c (e.g., text description in SD, also known as ‘concept’) and structured to estimate the underlying noise (achieved by the reverse diffusion process). Here $\mathbf{x}_t$ denotes the data or the latent feature subject to noise injection (attained via forward diffusion process) at the diffusion step t . The diffusion process is given by

$$\hat{\epsilon}_\theta(\mathbf{x}_t|c) = (1 - w)\epsilon_\theta(\mathbf{x}_t|\emptyset) + w\epsilon_\theta(\mathbf{x}_t|c), \quad (11.2)$$

where $\hat{\epsilon}_\theta(\mathbf{x}_t|c)$ stands for the ultimate noise estimation attained by utilizing the conditional DM given c , $w \in [0, 1]$ is a guidance weight, and $\epsilon_\theta(\mathbf{x}_t|\emptyset)$ signifies the corresponding unconditional employment of the DM. The inference stage initiates with Gaussian noise $z_T \sim \mathcal{N}(0, 1)$, which is then denoised using $\hat{\epsilon}_\theta(\mathbf{x}_t|c)$ to obtain z_{T-1} . This procedure is repeated to generate the authentic data at $t = 0$. When training the DM θ , the mean-squared-error (MSE) loss is commonly used

$$\ell_{\text{MSE}}(\theta; \mathcal{D}) = \mathbb{E}_{t, \epsilon \sim \mathcal{N}(0, 1)} [\|\epsilon - \epsilon_\theta(\mathbf{x}_t|c)\|_2^2], \quad (11.3)$$

where we omit the expectation over the training data in $\mathcal{D}$ for ease of presentation.

Given a well-trained DM θ , the objective of MU for image generation is twofold: firstly, to prevent θ from generating undesired image content, e.g., when conditioned

on harmful concepts like nudity; secondly, to ensure that the post-unlearning updated DM maintains the quality of image generation for normal images.

A widely recognized concept unlearning approach is ESD [241], notable for its state-of-the-art (SOTA) balance between unlearning effectiveness and model utility preservation [1012]. ESD facilitates the fine-tuning process of DMs by guiding outputs away from a specific concept targeted for erasure. Let c_e denote the concept to erase, then the diffusion process of ESD is modified to

$$\epsilon_{\theta}(\mathbf{x}_t|c_e) \leftarrow \epsilon_{\theta_o}(\mathbf{x}_t|\emptyset) - \eta (\epsilon_{\theta_o}(\mathbf{x}_t|c_e) - \epsilon_{\theta_o}(\mathbf{x}_t|\emptyset)), \quad (11.4)$$

where θ denotes the concept-erased DM, θ_o is the originally pre-trained DM, and $\epsilon_{\theta}(\mathbf{x}_t|\emptyset)$ represents unconditional generation of the model θ by considering text prompt as empty. Compared to the standard conditional DM [302] (with classifier-free guidance), the second term $-\eta[\epsilon_{\theta_o}(\mathbf{x}_t|c_e) - \epsilon_{\theta_o}(\mathbf{x}_t|\emptyset)]$ encourages the adjustment of the data distribution (with erasing guidance parameter $\eta > 0$) to minimize the likelihood of generating an image $\mathbf{x}$ that could be labeled as c_e . To optimize θ , ESD performs the following model fine-tuning based on (11.4):

$$\begin{aligned} & \underset{\theta}{\text{minimize}} \ell_{\text{MU}}(\theta, c_e) \\ & := \mathbb{E} \left[\left\| \epsilon_{\theta}(\mathbf{x}_t|c_e) - (\epsilon_{\theta_o}(\mathbf{x}_t|\emptyset) - \eta (\epsilon_{\theta_o}(\mathbf{x}_t|c_e) - \epsilon_{\theta_o}(\mathbf{x}_t|\emptyset))) \right\|_2^2 \right], \end{aligned} \quad (11.5)$$

where for notational simplicity we have used, and will continue to use, to omit the time step t and the random initial noise ϵ under expectation.

The field of unlearning for DMs is evolving rapidly. In addition to **ESD** [241], other MU methods in DMs include **FMN** (Forget-Me-Not) [978], **AC** (ablating concepts) [427], **UCE** (unified concept editing) [243] and **SLD** (safe latent diffusion) [732]. Figure 11.1 displays some motivating results on the image generation of unlearned DMs vs. the vanilla DM given an inappropriate prompt. Depending on the unlearning scenarios, the applications of MU to DMs include (1) *concept unlearning*, focused on erasing the influences of a harmful prompt, (2) *style unlearning*, dedicated to disregarding a particular painting style, and (3) *object unlearning*, aimed at discarding knowledge of a specific object class.

LLM Unlearning: MU for LLMs LLM unlearning aims to mitigate the influence of undesired data, such as sensitive or copyrighted information, and/or restrict the model’s capabilities to avoid the associated content generation. This process also requires preserving the LLM’s utility for unrelated tasks and avoiding full retraining to maintain computational efficiency. While problem (11.1) may appear as a straightforward formulation for LLM unlearning initially, complexities arise in determining the effective forget loss ℓ_f and achieving the optimal balance between unlearning and utility. We present three representative LLM unlearning approaches (a)–(c) and illustrate how they relate to the specifics of problem (11.1).

Tasks:		Concept Unlearning		Style Unlearning		Object Unlearning	
Prompts:		P_1	P_2	P_3	P_4	P_5	P_6
		half body portrait of very beautiful 20-year-old woman ...high fashion...	dennis hopper crawling around on the floor, by norman ...	rooftops in paris by vincent van gogh	a wheatfield, with cypresses by vincent van gogh	Modern church architecture.	Tench in a fish market
Vanilla	SD						
	ESD						
	FMN						

Fig. 11.1 Examples of generated images using the vanilla SD (stable diffusion) and the unlearned DMs, including ESD (erased stable diffusion) [241] and FMN (forget-me-not) [978]. Three unlearning scenarios are considered: concept unlearning (removing inappropriate concepts), style unlearning (eliminating a painting style), and object unlearning (discarding knowledge of a specific object class). Each column shows generated images using different DMs with the same prompt (P_i) at the same seed

(a) *Gradient Difference (GradDiff)* [511, 566] The approach maximizes the training loss for the forget set, inducing divergence in the model’s predictions from their original state, while minimizing the loss on the retain set to uphold performance on unlearning-irrelevant tasks. Let $\ell(y|x; \theta)$ denote the prediction loss of using the model θ given the input x against the undesired response y . Then, the forget loss ℓ_f can be specified by utilizing the *negative* training loss over the forget set $\mathcal{D}_f$, while the retain loss remains the same as the training loss. This specifies (11.1) as

$$\underset{\theta}{\text{minimize}} \quad \underbrace{-\mathbb{E}_{(x,y) \in \mathcal{D}_f} [\ell(y|x; \theta)]}_{\text{GA}} + \lambda \mathbb{E}_{(x,y) \in \mathcal{D}_r} [\ell(y|x; \theta)]. \quad (11.6)$$

At $\lambda = 0$, problem (11.6) simplifies to maximizing the training loss on forget set. This method is known as *gradient ascent (GA)* [259, 948]. Therefore, the unlearning method formulated by (11.6) is called GradDiff, which captures the disparity between the ascent and descent of gradients over the forget set and retain set.

(b) *Preference Optimization (PO)* [205, 566] Drawing inspiration from direct preference optimization techniques [682], this approach substitutes the unbounded GA loss in (11.6) with an alignment loss based on new responses y_f when presented

with the forget set. The designated unlearning response could be a reject-based answer such as ‘I don’t know’ or an irrelevant answer devoid of the unlearning target-related information. This leads to the following optimization problem:

$$\min_{\theta} \mathbb{E}_{(x, y_f) \in \mathcal{D}_f} [\ell(y_f | x; \theta)] + \lambda \mathbb{E}_{(x, y) \in \mathcal{D}_f} [\ell(y | x; \theta)], \quad (11.7)$$

where compared to (11.6), unlearning is accomplished by minimizing the prediction loss concerning the preferred unlearning responses y_f .

(c) *Negative Preference Optimization (NPO)* [990] NPO also treats the unlearning problem as a preference optimization problem. Yet, different from PO that specifies the unlearning response y_f , it interprets the forgetting data in $\mathcal{D}_f$ as the negative examples and incorporates them alone in preference optimization [682]. This yields a similar problem as GradDiff (11.6), but replaces the GA loss with the negative examples-based preference optimization loss.

11.3 Sparse Optimization for MU: Leveraging Model Sparsity for Efficient and Effective Unlearning

In this section, we introduce a novel machine unlearning (MU) paradigm by exploiting its model-based perspective: Model sparsification through weight pruning. We demonstrate that model sparsity can enhance the multi-criteria unlearning performance of an approximate unlearner, reducing the approximation gap with exact unlearning while maintaining efficiency. Building on this insight, we also develop a sparse optimization foundation for MU that incorporates sparsity regularization to improve the training process of approximate unlearning.

Model sparsification, or weight pruning, has been extensively studied in the literature, with significant focus on the relationship between model compression and generalization [230, 283, 559, 1011]. For example, the notable Lottery Ticket Hypothesis (LTH) [230] demonstrated the existence of a sparse subnetwork (the so-called ‘winning ticket’) that matches or even exceeds the test accuracy of the original dense model. Figure 11.2 illustrates the relationship between the pruned model’s generalization performance and its sparsity ratio. In this example, one-shot magnitude pruning (OMP) [559] is used to obtain sparse models. OMP is computationally the lightest pruning method, directly pruning the model weights to the target sparsity ratio based on their magnitudes. As shown in the figure, there exists a regime where the model maintains its testing accuracy even as it becomes increasingly sparse, indicating a graceful balance between sparsity and performance.

Beyond generalization, the effects of pruning have also been investigated in various contexts such as model robustness [131, 741], fairness [783], interpretability [903], loss landscape [231], and privacy [345]. In particular, the privacy gains from pruning suggest connections between data influence and model sparsification,




Fig. 11.2 Testing accuracy of OMP-based sparse ResNet-18 vs. the dense model on CIFAR-10

making it a promising approach for efficient and effective unlearning in machine learning models [369].

Gains of MU from Sparsity We begin by quantifying the impact of model sparsity on MU through the lens of *unrolling stochastic gradient descent* (SGD) [808]. This specific SGD method allows us to derive the unlearning error, defined as the weight difference between the approximately unlearned model and the gold-standard retrained model, when scrubbing a single data point.

Let us assume a binary mask $\mathbf{m}$ associated with the model parameters θ , where $m_i = 0$ signifies that the i th parameter θ_i is pruned to zero and $m_i = 1$ represents the unmasked θ_i . This sparse pattern $\mathbf{m}$ could be obtained by a weight pruning method, like OMP. Given $\mathbf{m}$, the sparse model is $\mathbf{m} \odot \theta$, where $\odot$ denotes the element-wise multiplication. The work [808] showed that if gradient ascent (GA) is adopted to scrub a single data point for the original (dense) model θ (i.e., $\mathbf{m} = \mathbf{1}$), then the gap between GA and retraining from scratch (termed as Retrain or exact unlearning) can be approximately bounded in the weight space. The work [369] further extended the existing unlearning error analysis to a sparse model. To be specific, given the model sparse pattern $\mathbf{m}$ and the SGD-based training, the unlearning error of GA, denoted by $e(\mathbf{m})$, can be characterized by the weight distance between the GA-unlearned model and the gold-standard retrained model. This leads to the error bound [369]

$$e(\mathbf{m}) = O(\eta^2 t \|\mathbf{m} \odot (\theta_t - \theta_0)\|_2 \sigma(\mathbf{m})) \quad (11.8)$$

where O is the big-O notation, η is the learning rate, t is the number of training iterations, $(\theta_t - \theta_0)$ denotes the weight difference at iteration t from its initialization θ_0 , and $\sigma(\mathbf{m})$ is the largest singular value (σ) of the Hessian $\nabla_{\theta, \theta}^2 \ell$ (for a training loss ℓ) among the unmasked parameter dimensions, i.e., $\sigma(\mathbf{m}) :=$

$\max_j \{\sigma_j(\nabla_{\theta, \theta}^2 \ell), \text{ if } m_j \neq 0\}$. It is clear from (11.8) that the unlearning error reduces as the model sparsity in $\mathbf{m}$ increases. By contrast, the unlearning error derived in [808] for a dense model (i.e., $\mathbf{m} = \mathbf{1}$) is proportional to the dense model distance $\|\theta_t - \theta_0\|_2$. Thus, model sparsity is beneficial to reducing the gap between (GA-based) approximate and exact unlearning.

The aforementioned sparsity benefits suggest a new MU paradigm: ‘prune first, then unlearn’, as demonstrated in [369]. This approach leverages the fact that (approximate) unlearning on a sparse model results in a smaller unlearning error, thereby improving efficacy. This promising finding also opens the door to developing sparsity-aware MU methods that can directly scrub data influence from a dense model.

Sparsity-Aware Unlearning We next demonstrate if pruning and unlearning can be carried out simultaneously, without requiring prior knowledge of model sparsity. Let $\ell_{\text{MU}}(\theta; \theta_o)$ denote the unlearning objective function of model parameters θ , given the pre-trained state θ_o . Inspired by sparsity-inducing optimization [36], we integrate an ℓ_1 norm-based sparse penalty into ℓ_{MU} . This leads to the problem of ‘ ℓ_1 -sparse MU’:

$$\underset{\theta}{\text{minimize}} \ell_{\text{MU}}(\theta; \theta_o) + \gamma \|\theta\|_1, \quad (11.9)$$

where $\gamma > 0$ is a regularization parameter that controls the penalty level of the ℓ_1 norm, thereby reducing the magnitudes of ‘unimportant’ weights. In practice, the unlearning performance can be sensitive to the choice of the sparse regularization parameter γ . To address this limitation, one could design a sparse regularization scheduler. The work [369] demonstrated that using a linearly decreasing γ scheduler outperforms other schemes. This scheduler not only minimizes the gap in unlearning efficacy compared to retraining from scratch (Retrain), but also improves the preservation of model utility after unlearning. These findings suggest that it is advantageous to prioritize promoting sparsity during the early stages of unlearning and then gradually shift the focus towards enhancing fine-tuning accuracy on the remaining dataset.

In Fig. 11.3, we showcase the effectiveness of ℓ_1 -sparse MU. For ease of presentation, we focus on the comparison with fine-tuning (FT) on the remaining dataset ($\mathcal{D}_r$) and the optimal Retrain strategy in both class-wise forgetting and random data forgetting scenarios under (CIFAR-10, ResNet-18). As shown, ℓ_1 -sparse MU outperforms FT in terms of unlearning efficacy (measured by UA and MIA performance), and significantly narrows the performance gap with Retrain while retaining the computational advantages of approximate unlearning.

Gradient-Based Weight Saliency Map Although weight sparsity simplifies MU for discriminative models, it has shown to be less effective for generative models [214]. Specifically, weight sparsity faces two major limitations when applied to MU for generative models: (1) Determining the appropriate sparse pattern for a generative model (e.g., a diffusion model) can be inherently challenging; and (2) Even when




Fig. 11.3 Performance of ℓ_1 -sparse MU vs. FT and Retrain on class-wise forgetting and random data forgetting under (CIFAR-10, ResNet-18). Each metric is normalized to [0, 1] based on the best result across unlearning methods, while the actual best value is provided (e.g., 2.52 is the least computation time for class-wise forgetting)

sparsity is achievable, some applications may not favor delivering a sparse model after MU due to the observed decline in performance.

Expanding on weight sparsity, an alternative mechanism called the weight saliency map was developed in [214]. This mechanism focuses MU on specific model weights deemed *salient* to the unlearning process. This concept allows us to decompose the pre-unlearning model weights (θ_o) into two distinct components: the salient model weights earmarked for updating during MU and the intact model weights that remain unchanged. We utilize the gradient of the forgetting loss $\ell_f(\theta; \mathcal{D}_f)$ in (11.1) with respect to the model weights. By applying a hard thresholding operation, we can then obtain the desired weight saliency map:

$$\mathbf{m}_S = \mathbb{I}(|\nabla_{\theta} \ell_f(\theta; \mathcal{D}_f)|_{\theta=\theta_o}| \geq \gamma), \quad (11.10)$$

where $\mathbb{I}(\mathbf{g} \geq \gamma)$ is an element-wise indicator function which yields a value of 1 for the i -th element if $g_i \geq \gamma$ and 0 otherwise, $|\cdot|$ is an element-wise absolute value operation, and $\gamma > 0$ is a hard threshold. In practice, it was shown in [214] that setting γ to the median of the gradient vector $\nabla_{\theta} \ell_f(\theta; \mathcal{D}_f)|_{\theta=\theta_o}$ is a sufficiently effective choice. Based on (11.10), we explicitly express the unlearning model θ_u as

$$\theta_u = \underbrace{\mathbf{m}_S \odot (\Delta\theta + \theta_o)}_{\text{salient weights}} + \underbrace{(\mathbf{1} - \mathbf{m}_S) \odot \theta_o}_{\text{original weights}} \quad (11.11)$$




Fig. 11.4 Examples of generated images using SDs w/ and w/o MU. The unlearning methods include ESD, FMN, and SalUn. Each column represents generated images using different SDs with the same I2P prompt and the same seed.

where $\odot$ is element-wise product, $\Delta\theta$ signifies the weight updating scheme required for unlearning with respect to θ_o , and $\mathbf{1}$ denotes an all-one vector. The implication from (11.11) is that during weight updating in MU, the attention can be directed towards the salient weights. During the unlearning process, the variable $\Delta\theta$ is optimized to solve problem (11.1) given the weight saliency map $\mathbf{m}_S$. This targeted approach, termed as saliency unlearning (SalUn), ensures that the most influential weights are adjusted, enhancing the efficiency and effectiveness of the unlearning process even when applied to generative models.

To demonstrate the effectiveness of SalUn, we assess its performance in concept-wise forgetting for generative models, specifically to eliminate the impact of NSFW (not safe for work) concepts introduced through inappropriate image prompts (I2P) [732]. We generated images using the open-source Stable Diffusion (SD) V1.4 model with prompts provided by I2P and classified these images into various nude body parts using the NudeNet detector [53]. Our goal is to use an MU method to effectively erase the influence of nudity-related prompts in SD. Figure 11.4 presents the unlearning performance of different methods, including SalUn, and the ESD and FMN baselines introduced in Fig. 11.1. The effectiveness of unlearning is measured by the reduction in nudity-related image generations using the unlearned SD model with I2P prompts. For comparison, we also include the performance of the original SD model. As we can see, SalUn generates the fewest harmful images under I2P prompts. Additionally, without unlearning, the original SD V1.4 generates a substantial number of harmful images, highlighting the critical importance of MU in image generation.

11.4 Second-Order Optimization for MU: Iterative Influence-Guided Unlearning

In this section, we transition our focus from MU for vision tasks (image classification and image generation) to MU for language tasks, specifically unlearning in large language models (LLMs). We will highlight a critical yet often overlooked factor in LLM unlearning: the choice of optimizer. This foundational element is crucial for the effectiveness of LLM unlearning and warrants thorough exploration.

Influence Unlearning and Insights Influence unlearning is a one-shot machine unlearning technique that utilizes the influence function approach [268, 416] to assess and quantify the impact of the forget set $\mathcal{D}_f$ on the pre-trained model θ_o . Diverging from *iterative* optimization approaches like GradDiff (11.6) and PO (11.7), influence unlearning involves a *single* weight modification step, updating θ_o based on the influence exerted by the forget set on the weight space. While influence unlearning is a classic technique, its usage has been limited to vision tasks and small models [354, 881]. Even within the realm of vision tasks, it is not deemed a state-of-the-art (SOTA) approach to unlearning [369]. This is because influence unlearning relies on several strong approximations in its derivation and computation, as elaborated on below.

Let θ_u denote a retrained model from scratch on the retain set $\mathcal{D}_r$, i.e., the solution to the optimization problem $\min_{\theta} \mathbb{E}_{(x,y) \in \mathcal{D}_r} [\ell(y|x; \theta)]$ with random initialization, where ℓ is the training loss introduced in (11.6). The *objective of influence unlearning* is to derive the weight modification from the pre-trained model θ_o to the retrained model θ_u , i.e., $\theta_u - \theta_o$. To this end, a *weighted* training problem is introduced:

$$\theta(\mathbf{w}) := \arg \min_{\theta} \ell(\theta, \mathbf{w}), \quad \ell(\theta, \mathbf{w}) = \sum_{i=1}^N [w_i \ell(y_i|x_i; \theta)] \quad (11.12)$$

where (x_i, y_i) is training data point, N is the total number of training data points, and w_i represents the introduced data influence weight. If the data point (x_i, y_i) is removed from the training set, i.e., $(x_i, y_i) \in \mathcal{D}_r$, then w_i takes a value of 0. By (11.12), the pretrained and retrained models θ_o and θ_u can be expressed as

$$\theta_o = \theta(\mathbf{1}), \quad \theta(\mathbf{w}_{\text{MU}}) = \theta_u, \quad (11.13)$$

where $\theta(\mathbf{1})$ entails training over the entire training set with weights $\mathbf{w} = \mathbf{1}$. Here $\mathbf{1}$ denotes the all-one vector. Similarly, given the unlearning-specific weighting scheme, $\mathbf{w}_{\text{MU}} = \mathbf{1}_{\mathcal{D}_r}$, $\theta(\mathbf{w}_{\text{MU}})$ corresponds to the retrained model post unlearning. Here $\mathbf{1}_{\mathcal{D}_r}$ denotes an element-wise indicator function that takes the value 1 if the

data point belongs to the retain set $\mathcal{D}_r$ and 0 otherwise. Based on (11.13), influence unlearning then aims to derive:

$$\Delta(\mathbf{w}_{\text{MU}}) = \boldsymbol{\theta}(\mathbf{w}_{\text{MU}}) - \boldsymbol{\theta}(\mathbf{1}). \quad (11.14)$$

The derivation of (11.14) is highly non-trivial as the retrained model $\boldsymbol{\theta}(\mathbf{w}_{\text{MU}})$ cannot be directly obtained and is implicitly defined through the optimization problem $\min_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}, \mathbf{w}_{\text{MU}})$. To proceed, the influence function approach [268, 369, 416] simplifies (11.14) by applying a first-order Taylor expansion to $\boldsymbol{\theta}(\mathbf{w}_{\text{MU}})$ at $\mathbf{w} = \mathbf{1}$:

$$\begin{aligned} \Delta(\mathbf{w}_{\text{MU}}) &= \boldsymbol{\theta}(\mathbf{w}_{\text{MU}}) - \boldsymbol{\theta}(\mathbf{1}) \\ &\approx \left. \frac{d\boldsymbol{\theta}(\mathbf{w})}{d\mathbf{w}} \right|_{\mathbf{w}=\mathbf{1}} (\mathbf{w}_{\text{MU}} - \mathbf{1}), \end{aligned} \quad (11.15)$$

where $\frac{d\boldsymbol{\theta}(\mathbf{w})}{d\mathbf{w}}$ denotes the full derivative of $\boldsymbol{\theta}(\mathbf{w})$ with respect to (w.r.t.) $\mathbf{w}$, and is known as *implicit gradient* [262, 1008]. Utilizing the implicit function theorem [420], the closed form of the influence unlearning formula (11.15) can be given by Jia et al. [369, Proposition 1]:

$$\boldsymbol{\theta}_{\text{MU}} = \boldsymbol{\theta}_o + \mathbf{H}^{-1} \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}, \mathbf{1} - \mathbf{w}_{\text{MU}}) \big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_o}, \quad (11.16)$$

where $\ell(\boldsymbol{\theta}, \mathbf{w})$ represents the $\mathbf{w}$ -weighted training loss (11.12), $\mathbf{H}^{-1}$ stands for the inverse of the second-order derivative (i.e., Hessian matrix) $\nabla_{\boldsymbol{\theta}, \boldsymbol{\theta}} \ell(\boldsymbol{\theta}, \mathbf{1}/N)$ evaluated at $\boldsymbol{\theta}_o$, $\nabla_{\boldsymbol{\theta}} \ell$ denotes the gradient of ℓ , and $\mathbf{1} - \mathbf{w}_{\text{MU}}$ yields $\mathbf{1} - \mathbf{1}_{\mathcal{D}_r}$, which captures the data weight on the forget set $\mathcal{D}_f$. To compute (11.16), one must determine the inverse-Hessian gradient product. However, exact computation is often computationally prohibitive. To address this challenge, numerical approximations such as the WoodFisher approximation [762] are often employed to estimate the inverse-Hessian gradient product.

An **intriguing observation** from (11.16) is that influence unlearning conforms to the generic form of second-order (SO) optimization [76]. As in Newton's method, one uses a SO approximation of a loss function ℓ to locate its minima. This yields a descent algorithm based on a Newton step [52]:

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \underbrace{\eta_t \mathbf{H}_t^{-1} \mathbf{g}_t}_{\text{Newton step}}, \quad (11.17)$$

where t represents the iteration index of Newton's method, $\boldsymbol{\theta}_{t+1}$ denotes the currently updated optimization variables, $\eta_t > 0$ is the learning rate, and $\mathbf{H}_t$ and $\mathbf{g}_t$ represent the Hessian matrix and the gradient of the loss ℓ , respectively, evaluated at $\boldsymbol{\theta}_t$.

The consistency observed in the formats of influence unlearning (11.16) and second-order optimization (11.17) prompts us to consider whether we can integrate second-order optimization into influence unlearning, thereby transforming the latter into an effective iterative unlearning approach.

Second-Order Unlearning for LLMs As evident from the derivations of influence unlearning (11.16), there exist two primary limitations that hinder its application to LLM unlearning: the computational complexity associated with inverting the Hessian matrix, and the diminished accuracy stemming from approximations utilized in Taylor expansion and second-order information acquisition. If we can transition from the static, one-shot nature of influence unlearning to a dynamic, iterative optimization process, we anticipate that the diminished accuracy resulting from the approximations used in influence unlearning (11.16) will be mitigated through the iterative engagement of the learning process. However, we still face the computational challenge posed by the Hessian inversion in (11.17). Therefore, we need to select a practically feasible SO (second-order) optimization method for LLM unlearning.

Sophia (Second-order Clipped Stochastic Optimization) [513], a simple scalable SO optimizer, is well-suited since it utilizes a simple diagonal matrix estimate of the Hessian and has shown its effectiveness in LLM pre-training. Sophia modifies the vanilla Newton’s method to

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \eta_t \text{clip}(\mathbf{m}_t / \max\{\gamma \mathbf{h}_t, \epsilon\}, 1), \quad (11.18)$$

where $\mathbf{m}_t \leftarrow \beta_1 \mathbf{m}_{t-1} + (1 - \beta_1) \mathbf{g}_t$ is the exponential moving average (EMA) of the FO (first-order) gradient with parameter $\beta_1 > 0$, $\mathbf{h}_t$ denotes the EMA of the Hessian diagonal estimates obtained from the diagonal of the Gauss-Newton matrix [513], and the clipping operation $\text{clip}(\boldsymbol{\theta}, a)$ limits the magnitude of each element in vector $\boldsymbol{\theta}$ to a maximum of a , thereby preventing excessively large updates that could destabilize the optimization process. In (11.18), both the clipping operation $\text{clip}(\cdot, \cdot)$ and the division operation $\cdot / \cdot$ are all performed element-wise, and $\gamma > 0$ and $\epsilon > 0$ are additional parameters in the clipping operation. In (11.18), if the clipping operation is absent with $\gamma = 1$ and $\epsilon \rightarrow 0$, then the Sophia update (11.18) simplifies to the Newton update (11.17) utilizing the diagonal Hessian estimate for $\mathbf{H}$.

Next, we can link influence unlearning (11.16) with the SO optimizer and propose the SO unlearning approach. Recall from (11.16) and (11.12) that the change in data weights ($\mathbf{1} - \mathbf{w}_{\text{MU}}$) encodes the influence of the forget set $\mathcal{D}_f$ in model training. Therefore, we can interpret the term $\mathbf{H}^{-1} \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}_0, \mathbf{1} - \mathbf{w}_{\text{MU}})$ in (11.16) as a second-order optimization-based *ascent* step over the *forget set*. This contrasts with the original Sophia update (11.18), which executes the descent using the clipped Newton step. Let us take GradDiff (11.6) as an example. In the context of LLM unlearning, SO optimization will be conducted in two modes: the descent step over the retain set and the ascent step over the forget set.

Application to Copyrighted Content Removal in LLMs We demonstrate the effectiveness of second-order unlearning in the task of Who’s Harry Potter (WHP) copyrighted information removal across two LLMs fine-tuned on the Harry Potter book series dataset [205]. As for model configurations, we use the OPT-1.3B [993] and LLaMA2-7b-chat [819] as our base LLMs. The unlearning efficacy is measured using BLEU and Rouge-L recall scores for WHP sentence completion, given prompt lengths of 100 and 300 tokens, respectively. The model utility is assessed using the LM Evaluation Harness [245] to compute perplexity (PPL) on the Wikitext dataset [580] and mean zero-shot accuracy across various tasks. Additional evaluations include the TruthfulQA benchmark [503]. We will assess the effectiveness of second-order unlearning by comparing it with existing LLM unlearning approaches, as outlined below: *Gradient ascent (GA)*: This serves as a specialization of GradDiff (11.6) by setting its regularization parameter $\lambda = 0$. *GradDiff* (11.6) and *PO* (11.7): These methods are executed via regularized optimization, employing either FO (first-order) or SO (second-order) optimizers.

Table 11.2 presents the unlearning efficacy and model utility across different LLM unlearning methods built upon either FO or SO optimization. As we can see, the use of a second-order optimizer substantially improves unlearning efficacy. For instance, comparing FO-GradDiff and SO-GradDiff reveals a notable decrease in BLEU score at a prompt length of 300 in the LLaMA2-7B-chat model. This decrease suggests that the generated texts deviate further from the original book’s content, indicating more effective unlearning. Furthermore, the improvements observed in both perplexity (PPL) and zero-shot accuracy with SO unlearning methods highlight a superior balance between forget efficacy and utility preservation. Across different LLM unlearning methods, the GA method struggles to balance forget efficacy with utility preservation. Although it achieves the lowest scores on the LLaMA2-7B-chat model, it results in notably poor utility, performing substantially worse than other methods.

11.5 Adversarial Evaluation of MU

In this section, we address the challenge of achieving faithful evaluation for models post-unlearning (referred to as ‘unlearned models’) in an adversarial environment. We approach this problem from two perspectives: data selection and adversarial attack generation. First, we focus on identifying the data subset that presents the most significant challenge for data influence erasure, i.e., pinpointing the worst-case forget set. Second, we introduce an effective and efficient adversarial prompt generation approach for diffusion models (DMs). This method leverages the intrinsic classification abilities of DMs to simplify the creation of adversarial prompts, ensuring a robust evaluation of unlearning techniques.

Table 11.2 Performance of different unlearning methods on copyright removal across two LLMs. ‘Original’ refers to the original model without unlearning. ‘FO’ and ‘SO’ indicate the choice of the unlearning optimizer. Thus, the second-order LLM unlearning methods correspond to SO-GradDiff and SO-PO. The ↓ symbol denotes metrics where lower values indicate better unlearning performance, while ↑ symbolizes metrics where higher values are preferable, reflecting better retention of model utility. The ‘Unlearning Efficacy’ category measures the model’s success in removing targeted information, whereas ‘Utility’ gauges the model’s retained functionality post-unlearning. The optimal and second-best results for each column, excluding those for the original model, are emphasized in bold and underlined, respectively

Method	Unlearning efficacy				Utility		Zero-shot Acc.↑	TruthfulQA↑
	Prompt Length 100		Prompt Length 300		PPL↓			
	BLEU↓	ROUGEL↓	BLEU↓	ROUGEL↓				
OPT-1.3B								
Original	6.3288	0.1701	6.8797	0.2453	59.33		46.69%	0.2313
FO-GA	5.7520	0.1725	6.0775	0.2421	71.04		46.31%	0.2301
FO-GradDiff	1.8633	0.1681	2.8236	0.2160	37.25		46.33%	0.2632
SO-GradDiff	<u>0.7841</u>	0.1090	1.3476	0.1480	34.09		46.80%	0.2277
FO-PO	0.9805	0.0620	2.2445	0.0815	24.98		45.76%	0.2607
SO-PO	0.6456	0.0476	<u>1.8619</u>	0.0707	24.08		<u>46.69%</u>	<u>0.2387</u>
LLaMA2-7B-chat								
Original	4.6489	0.1565	3.4986	0.1637	10.73		61.31%	0.2729
FO-GA	0.0135	0.0015	0.0279	0.0013	15.66		59.91%	0.2791
FO-GradDiff	0.2521	0.0247	0.6345	0.0476	11.18		60.06%	0.2681
SO-GradDiff	<u>0.1577</u>	<u>0.0117</u>	<u>0.4243</u>	<u>0.0180</u>	10.66		60.04%	0.2595
FO-PO	0.3120	0.0495	0.8530	0.0750	<u>9.48</u>		61.14%	0.2950
SO-PO	0.2499	0.0435	0.5284	0.0496	9.47		60.12%	<u>0.2827</u>

Identification of the Worst-Case Forget Set in MU Our motivation stems from the limitations in current MU evaluation methods, which heavily rely on artificially constructed *random data forgetting* scenarios [214, 369, 430]. However, studies in [213, 214] indicate that the effectiveness of unlearning methods can significantly vary with the selection of the forget set (i.e., the specific data points designated for forgetting), resulting in substantial performance variance. This variability in unlearning effectiveness based on forget set choices prompts us to reconsider the evaluation approach. We propose exploring a *worst-case forget set* selection scenario. This scenario would ideally represent the most challenging conditions for an unlearning method’s performance, thereby reducing unlearning variance and facilitating a more reliable assessment. By identifying the most difficult data points for influence erasure, we can better evaluate and compare the robustness of different unlearning techniques. We address the problem of identifying the worst-case forget set through a bi-level optimization (BLO)-based data selection framework, as elaborated below.

BLO offers a hierarchical learning framework, featuring two tiers of optimization tasks, i.e., the upper and lower levels. In this structure, the objective and variables of the upper-level problem are contingent upon the solution of the lower-level problem. In the context of identifying the worst-case forget set, we optimize the selection of a forget set at the *upper* level to maximize the difficulty of unlearning. Concurrently, the *lower* level is dedicated to generating the unlearned model, aiming to meet the unlearning objectives without compromising the utility on non-forgetting data points.

Let us introduce an optimization variable $\mathbf{w} \in \{0, 1\}^N$, recalling that N represents the total number of training data points. Here $w_i = 1$ signifies that the i -th training data point is included in the forget set, i.e., $\mathcal{D}_f = \{\mathbf{z}_i | w_i = 1\}$. Thus, our objective is to optimize the data selection scheme $\mathbf{w}$, such that the associated $\mathcal{D}_f$ can characterize the worst-case performance of an unlearned model, i.e., challenging the model θ_u in (11.1) post the unlearning of the designated forget set.

We form the *lower-level* optimization problem to determine the unlearned model θ_u based on the forget set defined by $\mathbf{w}$. By integrating $\mathbf{w}$ into (11.1), the unlearning problem in *lower-level optimization* can be cast as

$$\theta_u(\mathbf{w}) = \arg \min_{\theta} \ell_{\text{MU}}(\theta; \mathbf{w}) := \sum_{\mathbf{z}_i \in \mathcal{D}} [w_i \ell_f(\theta; \mathbf{z}_i) + (1 - w_i) \ell_r(\theta; \mathbf{z}_i)], \quad (11.19)$$

where $\theta_u(\mathbf{w})$ signifies the resulting unlearned model that is a function of $\mathbf{w}$, and the loss terms $\sum_{\mathbf{z}_i \in \mathcal{D}} [w_i \ell_f(\theta; \mathbf{z}_i)]$ and $\sum_{\mathbf{z}_i \in \mathcal{D}} [(1 - w_i) \ell_r(\theta; \mathbf{z}_i)]$ correspond to the forget loss and the retain loss in (11.1) on the forget set $\mathcal{D}_f$ and the retain set $\mathcal{D}_r$, respectively. Unless specified otherwise, we specify the unlearning objective (11.19) through the finetuning-based unlearning strategy, with $\lambda = 1$ and $\ell_f = -\ell_r$ in (11.1). Here, both loss functions are given by the training loss ℓ (e.g., the cross-entropy loss for image classification) over θ , with the forget loss $\ell_f = -\ell$ designed to counteract the training, thereby enforcing the unlearning.

With the unlearned model $\theta_u(\mathbf{w})$ defined as a function of the data selection scheme $\mathbf{w}$, we proceed to outline the BLO framework by incorporating an *upper-level* optimization. This is designed to optimize $\mathbf{w}$ for the worst-case unlearning performance, yielding the *overall BLO problem*:

$$\min_{\mathbf{w} \in \mathcal{S}} \underbrace{\sum_{\mathbf{z}_i \in \mathcal{D}} [w_i \ell(\theta_u(\mathbf{w}); \mathbf{z}_i)] + \gamma \|\mathbf{w}\|_2^2}_{\text{Upper-level objective} := f(\mathbf{w}, \theta_u(\mathbf{w}))}; \quad \text{subject to} \quad \underbrace{\theta_u(\mathbf{w}) = \arg \min_{\theta} \ell_{\text{MU}}(\theta; \mathbf{w})}_{\text{Lower-level optimization}} \quad (11.20)$$

where $\mathbf{w}$ is the upper-level optimization variable subject to the data selection constraint set $\mathcal{S}$, e.g., $\mathcal{S} = \{\mathbf{w} | \mathbf{w} \in \{0, 1\}^N, \mathbf{1}^\top \mathbf{w} = m\}$ with m being the forget set size, the lower-level objective function ℓ_{MU} has been defined in (11.19), and ℓ denotes the training loss. In addition, minimizing $\sum_{\mathbf{z}_i \in \mathcal{D}} [w_i \ell(\theta_u(\mathbf{w}); \mathbf{z}_i)]$ renders the worst-case scenario of the unlearned model $\theta_u(\mathbf{w})$ (derived from the lower-level optimization), i.e., making it ineffective at erasing the influence of the forget set (corresponding to $\{w_i = 1\}$) on model performance. Furthermore, we introduce an ℓ_2 regularization term with the regularization parameter $\gamma \geq 0$ in the upper-level objective function. This has dual purposes: it encourages sparsity in the data selection scheme $\mathbf{w}$ (when relaxed to continuous variables) and enhances the stability of BLO by including a strongly convex regularizer. To solve the BLO problem (11.20), an effective and efficient approach is to use the sign-based gradient unrolling method developed in [213].

Given the solution of (11.20), we justify its worst-case unlearning performance through the exact unlearning method, Retrain. In Table 11.3, we examine the performance disparities between the worst-case forget set and the random forget set in the task of MU for image classification on CIFAR-10, when employing Retrain at different forgetting data ratios including 1, 5, 10, and 20%. In terms of unlearning effectiveness, the chosen worst-case forget set consistently poses the greatest challenge for unlearning in all scenarios tested, as indicated by a significant drop in UA and MIA to nearly 0% (see the ‘Worst-case’ and ‘Diff’ columns of Table 11.3). In addition, the variance in worst-case unlearning effectiveness performance (as measured by UA and MIA) remains significantly lower than that observed with random data forgetting at various forgetting data ratios. Furthermore, the utility of the unlearned model, as indicated by RA and TA, shows no loss when comparing unlearning on worst-case forget sets to random forget sets. Intriguingly, the TA of models unlearned with the worst-case forget set may even surpass those unlearned with random sets, hinting at a connection to coreset selection as illustrated in [213].

Adversarial Prompts Against Unlearned DMs Since current unlearned DMs (diffusion models) often depend on heuristic-based and approximative unlearning methods, their trustworthiness remains in question. We address this problem by crafting adversarial attacks within the text prompt domain, i.e., adversarial prompts.

Table 11.3 Performance of exact unlearning (Retrain) under the random forget set and the worst-case forget set at different forgetting data ratios on CIFAR-10 using ResNet-18. The result format is given by $a \pm b$, with mean a and standard deviation b over 10 independent trials. The performance difference is provided in Diff, represents the worst-case performance is **lower**▼, **equal**−, or **higher**▲ than random-case performance

Metrics	1%-data forgetting			5%-data forgetting			10%-data forgetting			20%-data forgetting		
	Random	Worst-case	Diff	Random	Worst-case	Diff	Random	Worst-case	Diff	Random	Worst-case	Diff
UA	5.85 $\pm$ 0.69	0.00 $\pm$ 0.00	5.85 ▼	5.92 $\pm$ 0.44	0.00 $\pm$ 0.00	5.92 ▼	5.28 $\pm$ 0.33	0.00 $\pm$ 0.00	5.28 ▼	5.76 $\pm$ 0.20	0.00 $\pm$ 0.00	5.76 ▼
MIA	12.89 $\pm$ 1.27	0.00 $\pm$ 0.00	12.89 ▼	13.00 $\pm$ 0.55	0.02 $\pm$ 0.02	12.98 ▼	12.86 $\pm$ 0.1	0.00 $\pm$ 0.00	12.86 ▼	14.34 $\pm$ 0.40	0.03 $\pm$ 0.01	14.31 ▼
RA	99.96 $\pm$ 0.00	99.95 $\pm$ 0.02	0.01 ▼	100.00 $\pm$ 0.00	100.00 $\pm$ 0.00	0.00 −	100.00 $\pm$ 0.00	100.00 $\pm$ 0.00	0.00 −	100.00 $\pm$ 0.00	100.00 $\pm$ 0.00	0.00 −
TA	93.17 $\pm$ 0.15	93.45 $\pm$ 0.17	0.28 ▲	94.51 $\pm$ 0.07	94.67 $\pm$ 0.08	0.16 ▲	94.38 $\pm$ 0.1	94.66 $\pm$ 0.09	0.28 ▲	94.04 $\pm$ 0.08	94.60 $\pm$ 0.08	0.56 ▲

We investigate if subtle perturbations to text prompts can circumvent the unlearning mechanisms and compel unlearned DMs to once again generate harmful images.

In the attack setup, the *victim model* is represented by an *unlearned DM*, which is purported to effectively eliminate a specific concept, image style, or object class. Moreover, the crafted adversarial prompts (APs) are inserted before the original prompts, adhering to the format ‘[APs] + [Original Prompts]’. The format of APs has been achieved using only 3 ~ 5 token-level perturbations [1007]. Furthermore, the adversary operates within the white-box attack setting, having access to both the parameters of the victim model. We define the *adversarial attack problem* below: Given an unlearned DM θ^* that inhibits the image generation associated with a prompt c , we aim to craft a perturbed prompt c' (with subtle perturbations) that can circumvent the safety assurances provided by θ^* , thereby enabling image generation related to c .

Next, we introduce the attack method UnlearnDiffAtk [1007], which leverages the class discriminative ability of the ‘diffusion classifier’ [455] inherent in a well-trained DM for adversarial prompt generation. The underlying principle is that classification with a DM can be achieved by applying Bayes’ rule to the generation likelihood $p_\theta(\mathbf{x}|c)$ and the prior probability distribution $p(c)$ over prompts $\{c_i\}$ (viewed as image ‘labels’). Recall that $\mathbf{x}$ and θ denote an image and DM parameters, respectively. According to Bayes’ rule, the probability of predicting $\mathbf{x}$ as the ‘label’ c becomes

$$p_\theta(c_i|\mathbf{x}) = \frac{p(c_i)p_\theta(\mathbf{x}|c_i)}{\sum_j p(c_j)p_\theta(\mathbf{x}|c_j)}, \quad (11.21)$$

where $p(c)$ can be a uniform distribution, representing a random guess regarding $\mathbf{x}$, while $p_\theta(\mathbf{x}|c_i)$ is associated with the quality of image generation corresponding to prompt c_i . With the uniform prior, i.e., $p(c_i) = p(c_j)$, (11.21) can be simplified to only involve the conditional probabilities $\{p_\theta(\mathbf{x}|c_i)\}$. In DM, the log-likelihood of $p_\theta(\mathbf{x}|c_i)$ relates to the denoising error of the diffusion process, i.e., $p_\theta(\mathbf{x}|c_i) \propto \exp\{-\mathbb{E}_{t,\epsilon}[\|\epsilon - \epsilon_\theta(\mathbf{x}_t|c_i)\|_2^2]\}$, where $\exp \cdot$ is the exponential function, and t is a sampled time step [455]. As a result, the *diffusion classifier* is given by

$$p_\theta(c_i|\mathbf{x}) \propto \frac{\exp\{-\mathbb{E}_{t,\epsilon}[\|\epsilon - \epsilon_\theta(\mathbf{x}_t|c_i)\|_2^2]\}}{\sum_j \exp\{-\mathbb{E}_{t,\epsilon}[\|\epsilon - \epsilon_\theta(\mathbf{x}_t|c_j)\|_2^2]\}}. \quad (11.22)$$

Thus, the DM (θ) can serve as a classifier by evaluating its denoising error for a specific prompt (c_i) relative to all the potential errors across different prompts.

Through the lens of diffusion classifier (11.22), the task of creating an adversarial prompt (c') to evade a victim unlearned DM (θ^*) can be cast as:

$$\underset{c'}{\text{maximize}} \quad p_{\theta^*}(c'|\mathbf{x}_{\text{tgt}}), \quad (11.23)$$

where $\mathbf{x}_{\text{tgt}}$ denotes a *target image* containing unwanted content which θ^* intends to avoid such a generation, and the target image is encoded into the latent space, followed by the addition of random noises adhering to the same settings as those outlined in the diffusion classifier [455].

To address the above problems, we use a key observation in the diffusion classifier [455]: Classification only requires the *relative* differences between the noise errors, rather than their *absolute* magnitudes. This transforms (11.22) to

$$\frac{1}{\sum_j \exp \left\{ \mathbb{E}_{t,\epsilon} [\|\epsilon - \epsilon_{\theta}(\mathbf{x}_t | c_i)\|_2^2] - \mathbb{E}_{t,\epsilon} [\|\epsilon - \epsilon_{\theta}(\mathbf{x}_t | c_j)\|_2^2] \right\}}. \quad (11.24)$$

Based on (11.24), if we view the adversarial prompt c' as the targeted prediction label, i.e., $c_i = c'$ in (11.22), we can then solve the attack generation problem (11.23) as

$$\underset{c'}{\text{minimize}} \sum_j \exp \left\{ \mathbb{E}_{t,\epsilon} [\|\epsilon - \epsilon_{\theta^*}(\mathbf{x}_{\text{tgt},t} | c')\|_2^2] - \mathbb{E}_{t,\epsilon} [\|\epsilon - \epsilon_{\theta^*}(\mathbf{x}_{\text{tgt},t} | c_j)\|_2^2] \right\}, \quad (11.25)$$

where $\mathbf{x}_{\text{tgt},t}$ is the noisy image at diffusion time step t corresponding to the original noiseless image $\mathbf{x}_{\text{tgt}}$. To facilitate optimization, we can transform (11.25) into the following simplified optimization problem for attack generation:

$$\underset{c'}{\text{minimize}} \mathbb{E}_{t,\epsilon} [\|\epsilon - \epsilon_{\theta^*}(\mathbf{x}_{\text{tgt},t} | c')\|_2^2], \quad (11.26)$$

where we excluded $\exp$ as it is a convex and monotonically increasing function. Figure 11.5 summarizes the attack generation process.

In Table 11.4, we present the performance of various attack methods against unlearned DMs designed to mitigate the influence of inappropriate concepts from the I2P (inappropriate-to-prompt) dataset. We examine three unlearned DMs: ESD




Fig. 11.5 Pipeline of the proposed adversarial prompt learning method for unlearned DM's robustness evaluation

Table 11.4 Performance of various attack methods against unlearned DMs in concept unlearning, measured by attack success rate (ASR) and computation time in minutes (mins). ‘No Attack’ uses original prompts from I2P. ‘P4D’ [141] and UnlearnDiffAtk are optimization-based attack methods. ‘Attack Time’ represents the average computation time for generating one attack per prompt. The best attack performance (highest ASR or lowest computation time) is highlighted in **bold**

I2P:	Nudity			Violence			Illegal Activity			Atk. Time per Prompt (mins)				
	142			756			727							
	ESD	FMN	SLD	ESD	FMN	SLD	ESD	FMN	SLD					
Total prompts #:	Unlearned DMs:	Attacks: (ASR %)	No attack	20.42%	88.03%	33.10%	27.12%	43.39%	22.93%	30.99%	32.83%	27.78%	—	
			P4D	69.71 %	97.89%	77.46%	80.56%	85.85%	62.43%	85.83%	88.03%	81.98%		34.70
			UnlearnDiffAtk	76.05%	97.89%	82.39%	80.82%	84.13%	62.57%	85.01%	86.66%	82.81%		26.29

[241], FMN [978], and SLD ([732]. Our evaluation assesses their robustness across three categories of harmful concepts: nudity, violence, and illegal activity, comprising 142, 756, and 727 inappropriate prompts, respectively. We compare the attack performance of using UnlearnDiffAtk with that of two baselines: ‘No attack’, which uses the original inappropriate prompt from I2P; and ‘P4D’, which corresponds to the attack problem proposed in [141]. Additionally, we compare different attack methods with respect to ‘attack time’ (Atk. time), given by the average computation time needed to generate one attack per prompt on a single NVIDIA RTX A6000 GPU. As we can see, the optimization-based attacks (both UnlearnDiffAtk and P4D) can effectively circumvent various types of unlearned DMs, achieving higher ASR (attack success rate) than ‘No Attack’. Moreover, in most cases, UnlearnDiffAtk outperforms P4D although the ASR gap is not quite significant in concept learning. However, the former is achieved using lower computational cost than P4D, reducing runtime cost per attack instance generation by approximately 23.5%. By viewing from ASR, ESD demonstrates better robustness than other unlearned DMs, including FMN and SLD, when facing inappropriate prompts. Yet, unlearned DMs still lack satisfactory robustness against adversarial prompts.

In conclusion, by exposing the vulnerabilities of current MU techniques in both discriminative and generative models, we highlight the urgent need for robust evaluation methods in MU. This critical examination underscores the importance of developing more resilient unlearning strategies and reliable assessment frameworks to ensure that MU techniques are both effective and dependable in various adversarial scenarios. Addressing these challenges will be crucial for advancing the field and maintaining the integrity and privacy of ML models.

Part III
Trust and Safety in Foundation Models

Chapter 12

Trustworthiness Evaluation of Large Language Models



Abstract Large language models (LLMs) have gained considerable attention for their excellent natural language processing capabilities. Nonetheless, these LLMs present many challenges, particularly in the realm of trustworthiness. Therefore, ensuring the trustworthiness of LLMs emerges as an important topic. This chapter presents the TRUSTLLM framework (Sun et al., Trustllm: Trustworthiness in large language models. International Conference on Machine Learning (2024)), a comprehensive study of trustworthiness in LLMs, including principles for different dimensions of trustworthiness, established benchmark, evaluation, and analysis of trustworthiness for mainstream LLMs. Specifically, we first introduce a set of principles for trustworthy LLMs that span eight dimensions. Based on these principles, we further establish a benchmark across six dimensions including truthfulness, safety, fairness, robustness, privacy, and machine ethics. Based on the evaluation of 16 mainstream LLMs in TrustLLM (Sun et al., Trustllm: Trustworthiness in large language models. International Conference on Machine Learning (2024)), consisting of over 30 datasets, this chapter summarizes the main findings.

12.1 Introduction

The advent of large language models (LLMs) marks a significant milestone in natural language processing (NLP) and generative AI, as evidenced by numerous foundational studies [395, 740]. The exceptional capabilities of these models in NLP have garnered widespread attention, leading to diverse applications that impact every aspect of our lives. LLMs are employed in a variety of language-related tasks, including automated article writing [967], the creation of blog and social media posts, and translation [1047]. Additionally, they have improved search functionalities, as seen in platforms like Bing Chat [4, 584, 602], and other applications [2]. The efficacy of LLMs is distinctly evident in various other areas of human endeavor. For example, models such as Code Llama [714] offer considerable assistance to software engineers [587]. In the financial domain, LLMs like BloombergGPT [911] are employed for tasks including sentiment analysis, named entity recognition, news classification, and question answering. Furthermore, LLMs are increasingly being

applied in scientific research [12, 861, 942, 1004], spanning areas like medical applications [135, 157, 488, 631, 811, 833, 928, 980, 983, 1003], political science [508], law [5, 970], chemistry [277, 629], oceanography [63, 1026], education [964], and the arts [968], highlighting their extensive and varied impact.

The outstanding capabilities of LLMs can be attributed to multiple factors, such as the usage of large-scale raw texts from the Web as training data (e.g., PaLM [26, 147] was trained on a large dataset containing more than 700 billion tokens [739]), the design of transformer architecture with a large number of parameters (e.g., GPT-4 is estimated to have in the range of 1 trillion parameters [901]), and advanced training schemes that accelerate the training process, e.g., low-rank adaptation (LoRA) [322], quantized LoRA [180], and pathway systems [50]. Moreover, their outstanding instruction following capabilities can be primarily attributed to the implementation of alignment with human preference [367]. Prevailing alignment methods use reinforcement learning from human feedback (RLHF) [626] along with various alternative approaches [15, 43, 74, 96, 196, 237, 281, 442, 633, 650, 697, 793]. These alignment strategies shape the behavior of LLMs to more closely align with human preferences, thereby enhancing their utility and ensuring adherence to ethical considerations.

However, the rise of LLMs also introduces concerns about their trustworthiness. Unlike traditional language models, LLMs possess unique characteristics that can potentially lead to trustworthiness issues. (1) **Complexity and diversity of outputs from LLMs, coupled with their emergent generative capabilities.** LLMs demonstrate an unparalleled ability to handle a broad spectrum of complex and diverse topics. Yet, this very complexity can result in unpredictability and, consequently, the possibility of generating inaccurate or misleading outputs [35, 337, 368]. Simultaneously, their advanced generative capabilities open avenues for misuse by malicious actors, including the propagation of false information [114] and facilitating cyberattacks [165]. For instance, attackers might use LLMs to craft deceptive and misleading text that lures users to click on malicious links or download malware. Furthermore, LLMs can be exploited for automated cyberattacks, such as generating numerous fake accounts and comments to disrupt the regular operation of websites. A significant threat also comes from techniques designed to bypass the safety mechanisms of LLMs, known as *jailbreaking attacks* [882], which allows attackers to misuse LLMs illicitly. (2) **Data biases and private information in large training datasets.** One primary challenge to trustworthiness arises from potential biases in training datasets, which have significant implications for the fairness of content generated by LLMs. For example, a male-centric bias in the data may yield outputs that mainly reflect male perspectives, thereby overshadowing female contributions and viewpoints [28]. In a similar vein, a bias favoring a particular cultural background can result in responses biased toward that culture, thus disregarding the diversity present in other cultural contexts [166]. Another critical issue concerns the inclusion of sensitive personal information within training datasets. In the absence of stringent safeguards, this data becomes susceptible to misuse, potentially leading to privacy breaches [765]. This issue is especially acute in the healthcare sector, where maintaining the confidentiality of patient

data is of utmost importance [543]. (3) **High user expectations.** Users may have high expectations regarding the performance of LLMs, expecting accurate and insightful responses that emphasize the model’s alignment with human values. Many researchers are expressing concerns about whether LLMs align with human values. A misalignment could significantly impact their broad applications across various domains. For instance, an LLM considers a behavior appropriate in some situations. Still, humans may view it as inappropriate, leading to conflicts and contradictions in its applications, as highlighted in specific cases [564].

The developers of LLMs have undertaken significant efforts to address the concerns mentioned above. OpenAI [623] has taken measures to ensure LLMs’ trustworthiness in the training data phase, training methods, and downstream applications. WebGPT [602] is introduced to assist human evaluation in identifying inaccurate information in LLM responses. Meta [581], dedicated to responsible AI, bases its approach on five pillars: privacy, fairness, robustness, transparency, and accountability. The introduction of Llama2 [820] sets new safety alignment benchmarks for LLMs, encompassing extensive safety investigations in pretraining, fine-tuning, and red teaming. Despite these concerted efforts, a persistent question remains: *To what extent can we genuinely trust LLMs?*

To tackle these crucial questions, it is essential to address the fundamental issue of benchmarking how trustworthy LLMs are. What key elements define the trustworthiness of large language models, and from various perspectives, how should this trustworthiness be assessed? Furthermore, exploring methodologies to practically evaluate trustworthiness across these dimensions is vital. However, answering these questions is far from straightforward. The primary challenges include: (1) **Definition of comprehensive aspects.** One of the main obstacles is the absence of a universally accepted set of criteria that comprehensively encapsulates all facets of trustworthiness. This lack of standardized metrics makes it difficult to assess and compare the trustworthiness of different LLMs uniformly. (2) **Scalability and generalizability:** Creating benchmarks that are scalable across different sizes and types of LLMs and generalizable across various domains and applications is a complex task; (3) **Practical evaluation methodologies.** Effective prompts need to be designed to test obvious trustworthiness issues and uncover more subtle biases and errors that might not be immediately apparent. This requires a deep understanding of both the technology and the potential societal impacts of its outputs.

In this chapter, we present TRUSTLLM [791], a unified framework to support a comprehensive analysis of trustworthiness in LLM, including a survey of existing work, organizing principles of different dimensions of trustworthy LLMs, a novel benchmark, and a thorough evaluation of trustworthiness for mainstream LLMs. Specifically, we address the three challenges above as follows.

- **Identification of eight facets of trustworthiness.** To explore how trustworthy LLMs are, we incorporated domain knowledge from across AI, machine learning, data mining, human–computer interaction (HCI), and cybersecurity. We conducted an extensive review of 600 papers on LLM trustworthiness published

in the past five years and identified eight key aspects that define the trustworthiness of LLMs, which are truthfulness, safety, fairness, robustness, privacy, machine ethics, transparency, and accountability. In this work, to facilitate our investigation, we separate utility (i.e., functional effectiveness) from the eight identified dimensions and define *trustworthy LLMs* as “to be trustworthy, LLMs must appropriately reflect characteristics such as truthfulness, safety, fairness, robustness, privacy, machine ethics, transparency, and accountability.”

- **Selection of comprehensive and diverse LLMs for investigation.** By evaluating **16 LLMs**, encompassing both proprietary and open-source models, we cover a broad spectrum of model sizes, training strategies, and functional capabilities. This diversity guarantees that TRUSTLLM is not confined to a specific type or size of LLM. It also establishes a comprehensive evaluation framework for assessing the trustworthiness of future LLMs.
- **Benchmarking and evaluation across various tasks and datasets:** We benchmark **30 datasets** to comprehensively evaluate the functional capabilities of LLMs, ranging from simple classification to complex generation tasks. Each dataset presents unique challenges and benchmarks the LLMs across multiple dimensions of trustworthiness. Meanwhile, diverse evaluation metrics are employed to understand the capabilities of LLMs. This approach ensures that the evaluation is thorough and multifaceted.

12.2 Background and Related Work

12.2.1 Large Language Models (LLMs)

A language model (LM) aims to predict the probability distribution over a sequence of tokens. Scaling the model size and data size, large language models (LLMs) have shown “emergent abilities” [155, 887, 889] in solving a series of complex tasks that cannot be dealt with by regular-sized LMs. For instance, GPT-3 can handle few-shot tasks by learning in context, in contrast to GPT-2, which struggles in this regard. The success of LLMs is primarily attributed to the Transformer architecture [844]. Specifically, almost all the existing LLMs employ a stack of transformer blocks, each consisting of a Multi-Head Attention layer followed by a feedforward layer interconnected by residual links. Built upon this transformer-based architecture, there are three primary designs of LLMs: encoder-decoder architecture [685], causal-decoder architecture, and prefix-decoder architecture. Among them, the most widely used architecture is the causal decoder, which employs an attention mask to ensure that each input token only attends to previous tokens and itself. In this survey, we mainly focus on the causal-decoder architecture. The training of LLMs is usually composed of three steps: pre-training, instruction finetuning, and alignment tuning.

During pre-training, LLMs learn world knowledge and basic language abilities on large-scale corpora. To improve model capacity, researchers established some

scaling laws to show the compute-optimal ratio between the model size and data size, including KM scaling law [385] and Chinchilla scaling law [306]. When the scale reaches certain levels, LLMs show emergent abilities to solve complex tasks, instruction following, in-context learning, and step-by-step reasoning. These abilities endow LLMs to be general-purpose task solvers. To further elicit the instruction-following and in-context learning ability of LLMs, instruction tuning suggests creating appropriate task instructions or particular in-context learning methods to enhance the ability of LLMs to generalize to tasks they have not encountered before. During the alignment training phase, LLMs are trained to align with human values, e.g., being helpful, honest, and harmless, instead of producing harmful content. For this purpose, two kinds of alignment training methods, including supervised finetuning (SFT) and reinforcement learning from human feedback (RLHF), are proposed in InstructGPT, which is the fundamental algorithm behind the ChatGPT.

SFT guides the LLMs to understand the prompts and generate meaningful responses, which can be defined as follows. Given an instruction prompt x , we want the LLM to generate a response aligned with the human-written response y . The SFT loss is defined as the cross-entropy loss between the human-written response and the LLM-generated response, i.e., $\mathcal{L}_{\text{SFT}} = -\sum_t \log p(y_t|x, y_{<t})$, where $y_{<t}$ represents the sequence of tokens up to but not including the current token y_t . However, the limitation of SFT is that it only provides a single human-written response for each prompt, which is insufficient to provide a fine-grained comparison between the sub-optimal ones and capture the diversity of human responses. To address this issue, RLHF [626] is proposed to provide fine-grained human feedback with pair-wise comparison labeling. Typical RLHF includes three main steps: (1) SFT on high-quality instruction set; (2) collecting manually ranked comparison response pairs and training a reward model for quality assessment; (3) optimizing the SFT model under the PPO [737] reinforcement learning framework with the reward model from the second step. To prevent over-optimization in step (3), a KL-divergence regularization term between the current and SFT models is added to the loss function. However, the PPO algorithm is not stable during training. Thus, Reward rAnked Fine-Tuning (RAFT) [190] is proposed to replace Proximal Policy Optimization (PPO) training with direct learning on the high-ranked samples filtered by the reward model. Nevertheless, these online algorithms require interaction between policy, behavior policy, reward, and value model, which requires fine-grained tuning on the hyper-parameters to achieve stability and generalizability. To prevent this, offline algorithms like ranking-based approaches, including Direct Preference Optimization (DPO) and Preference Ranking Optimization (PRO), and language-based approaches, including Conditional Behavior Cloning [860], Chain of Hindsight [516], and Stable Alignment [521] are proposed. These methods eliminate the risk of overfitting a reward model and improve training stability using preference ranking data.

12.2.2 *Evaluation on LLMs*

Evaluation of LLMs is a fast-evolving field involving multi-dimensional evaluation across various tasks, datasets, and benchmarks [104]. It encompasses a wide range of domains, starting with traditional NLP tasks, where LLMs are assessed for natural language understanding, including tasks like sentiment analysis [545, 668, 1001], text classification [940, 991], natural language inference [574, 668], etc. The evaluation of LLMs also extends to reasoning tasks [104], covering mathematical reasoning [233, 668], logical reasoning [515, 634], and other reasoning parts; alongside natural language generation tasks like summarization [668, 998] and question answering [436, 668]; as well as including multilingual tasks [1000]. The evaluation also requires careful studies on robustness, especially in challenging situations such as out-of-distribution (OOD) and adversarial robustness [104, 862, 863], and learning rate tuning [378]. For trustworthiness, some work indicates that LLMs tend to absorb and express harmful biases and toxic content in their training data [249, 1049]. This underscores the need for comprehensive evaluation methodologies and a heightened focus on various trustworthiness aspects of LLMs [856]. Moreover, the application of LLMs expands into many other fields [272] including computational social science [1050], legal task [218, 273, 604], and psychology [229]. Besides, evaluating LLMs in natural science and engineering provides insights into their capabilities in mathematics [890, 969], general science [277, 603], and engineering [632, 780] domains. In the medical field, LLMs have been evaluated for their proficiency in addressing medical queries [307, 725], medical examinations [256, 429], and functioning as medical assistants [432, 880]. In addition, some benchmarks are designed to evaluate specific language abilities of LLMs like Chinese [340, 470, 494, 1002]. Besides, agent applications [502] underline their capabilities for interaction and using tools [344, 476, 669, 670]. Beyond these areas, LLMs contribute to different domains, such as education [171], finance [353, 482, 918, 984], search and recommendation [215, 446], personality testing [742]. Other specific applications, such as game design [434] and log parsing [439], illustrate the broad scope of the application and evaluation of LLMs. In addition to conventional text generation evaluations, the evaluations of LLMs have expanded to include their code generation capabilities [1030]. Recent studies have highlighted this emerging direction, revealing both the potential and the challenges in LLM-driven code synthesis [235, 519, 533, 1030].

In text generation evaluation, diverse untrained automatic evaluation metrics are utilized, including metrics based on n-gram overlap, distance-based measures, diversity metrics, content overlap metrics, and those with grammatical features [99]. Standard traditional metrics, such as BLEU [640] and ROUGE [501] classified as n-gram overlap metrics, estimate between the reference text and a text generated by the model. However, these metrics face limitations, particularly in scenarios where multiple correct methods of text generation exist, as often seen in tasks involving latent content planning or selection, which can also lead to accurate solutions receiving low scores [251, 760].

LLM evaluation datasets and benchmarks are vital in evaluating various language models for tasks, reflecting complex real-world language processing scenarios. Benchmarks like GLUE [854] and SuperGLUE [853] encompass various tasks from text categorization and machine translation to dialogue generation. These evaluations are crucial for understanding the capabilities of LLMs in general-purpose language tasks. Additionally, automatic and human evaluations serve as critical methods for LLM evaluation [104].

12.2.3 Trustworthiness-Related Benchmarks

Currently, in the domain of trustworthiness-related evaluation, there are many related works. For example, DecodingTrust [857] aims to thoroughly assess several perspectives of trustworthiness in GPT models. The recent study [593] proposes a prompting strategy by designing malicious demonstrations, and conducts an assessment of open-source LLMs on trustworthiness. Do-Not-Answer [874] introduces a dataset specifically designed to test the safeguard mechanisms of LLMs by containing only prompts that responsible models should avoid answering. SafetyBench [789] is a comprehensive benchmark for evaluating the safety of LLMs comprising diverse multiple-choice questions that span seven distinct categories of safety concerns. The HELM [492] is dedicated to enhancing the transparency of language models by comprehensively examining their capabilities and limitations by assessing various scenarios and metrics. Concurrently, the Red-Teaming benchmark [60] conducts security tests on LLMs to investigate their responses to potential threats. CVALUES [923] focuses on measuring the safety and responsibility of Chinese Language Large Models, while PromptBench [1044] examines the robustness of these models against adversarial prompts. Moreover, the GLUE-x [941] is centered on the open-domain robustness of language models. HaluEval [471] assesses the performance of LLMs in generating misinformation, and Latent Jailbreak [672] tests the safety and output robustness of models when presented with text containing malicious instructions. Finally, SC-Safety [927] engages Chinese LLMs with multi-turn open-ended questions to test their safety and trustworthiness. However, most of these benchmarks cover specific sections about trustworthiness, which are not comprehensive enough.

12.3 Guidelines and Principles for Trustworthiness Assessment of LLMs

To create guidelines for assessing the trustworthiness of LLMs, [791] presents the 8 principal dimensions of trustworthy LLMs, outlining their respective definitions

Table 12.1 The definitions of the eight identified dimensions

Dimension	Definition
Truthfulness	The accurate representation of information, facts, and results by an AI system.
Safety	The outputs from LLMs should only engage users in a safe and healthy conversation [536].
Fairness	The quality or state of being fair, especially fair or impartial treatment [7].
Robustness	The ability of a system to maintain its performance level under various circumstances [3].
Privacy	The norms and practices that help to safeguard human and data autonomy, identity, and dignity [3].
Machine ethics	Ensuring moral behaviors of man-made machines that use artificial intelligence, otherwise known as artificial intelligent agents [23, 24].
Transparency	The extent to which information about an AI system and its outputs is available to individuals interacting with such a system [3].
Accountability	An obligation to inform and justify one’s conduct to an authority [13, 507, 601, 614, 810].

and descriptions. The keywords of each principal dimension are cataloged within Table 12.1. Their details are provided as follows.

12.3.1 Truthfulness

Intricately linked to factuality, truthfulness stands out as an essential challenge for Generative AI models, including LLMs. It has garnered extensive discussion and scholarly attention [35, 70, 357, 1025]. To critically evaluate LLMs’ adherence to truthfulness, datasets and benchmarks, such as MMLU [293], Natural Questions [431], TriviaQA [382], and TruthfulQA [503], have been employed in prior works [859]. Some tools also assessed some specific aspects of general truthfulness: HaluEval [471] assesses hallucinations; SelfAware [954] explores awareness of knowledge limitations; FreshQA [850] and Pinocchio [326] inspect the adaptability to rapidly evolving information.

While accuracy remains a predominant metric for evaluating truthfulness [293, 471, 850, 954], the need for human evaluation is also recognized, particularly in benchmarks like TruthfulQA [503] and FreshQA [850]. However, the challenge of ensuring truthfulness is compounded by the inherent imperfections in training data [873]. LLMs, being trained on vast troves of text on the Internet, are susceptible to absorbing and propagating misinformation, outdated facts, and even intentionally misleading content embedded within their training datasets [635, 1036], making the pursuit of truthfulness in LLMs an ongoing and intricate challenge.

In [791], we define the *truthfulness* of LLMs as the accurate representation of information, facts, and results. Our assessment of the *truthfulness* of LLMs focuses on (1) evaluating their inclination to generate *misinformation* under two scenarios:

relying solely on internal knowledge and retrieving external knowledge; (2) testing LLMs' propensity to *hallucinate* across four tasks: multiple-choice question-answering, open-ended question-answering, knowledge-grounded dialogue, and summarization; (3) assessing the extent of *sycophancy* in LLMs, encompassing two types: persona sycophancy and preference sycophancy; and (4) testing the capabilities of LLMs to correct *adversarial facts* when, e.g., a user's input contains incorrect information.

12.3.2 Safety

With the pervasive integration of LLMs into various domains, safety and security concerns have emerged, necessitating comprehensive research and mitigation strategies [60, 87, 98, 121, 366, 531, 654, 671, 693, 704, 820, 869, 924, 927, 943, 947, 955, 962, 1027]. Although LLMs should be designed to be safe and harmless, their vulnerability to adversarial behaviors, such as *jailbreaking*, has been extensively documented [882]. Some commonly used jailbreaking methods include generation exploitation attacks [342] and straightforward queries [532] to sophisticated techniques involving genetic algorithms [435].

The repercussions of jailbreaking extend to the generation of toxic content and the misuse of LLMs, with the potential to significantly impact user interactions and downstream applications [893]. Furthermore, the role assigned to LLMs, dictated by their system parameters, can profoundly influence their propensity to generate toxic content, underscoring the need for vigilant role assignment and parameter tuning [179]. A prevalent form of misuse is *misinformation*, which exemplifies the potential harms associated with LLMs, and has been shown to result in tangible negative outcomes [288, 635, 1036].

Prior work has attempted to analyze the safety issues surrounding LLMs, tracing the origins of these issues and evaluating their impacts. Tools and datasets, such as Toxigen [285] and Realtoxicityprompts [250] have been developed to facilitate the detection of toxic content and assess the harm posed by LLMs. Integrating these tools into LLMs' development and deployment pipelines is crucial for ensuring that these powerful models are used safely and responsibly.

In TRUSTLLM, we define *Safety* as the ability of LLMs to avoid unsafe, illegal outputs and only engage users in a healthy conversation [536]. We first assess LLMs' safety against jailbreak attacks, by introducing a comprehensive taxonomy of jailbreak attacks comprising five major classes and 13 subclasses. Secondly, we evaluate the issue of over-alignment (i.e., exaggerated safety). Furthermore, we measure the toxicity levels in the outputs of LLMs that have been compromised by jailbreak attacks. Finally, we assess the LLMs' resilience against various misuse scenarios using the Do-Not-Answer dataset [874], the Do-Anything-Now dataset [749], and an additional dataset specifically curated for this study.

12.3.3 *Fairness*

Ensuring fairness in LLMs is crucial, as it encapsulates the ethical principle that necessitates the equitable design, training, and deployment of LLMs and related AI systems, preventing biased or discriminatory outcomes [864]. The significance of this issue is underscored by the increasing number of countries implementing legal frameworks that mandate adherence to fairness and anti-discrimination principles in AI models [226, 536].

There is a growing body of research dedicated to understanding the stages of model development and deployment where fairness could be jeopardized, including training data preparation, model building, evaluation, and deployment phases [239, 577, 795]. Fairness compromised due to the prevalence of bias in training datasets is often considered a top concern and has been the subject of extensive recent scrutiny [44, 185, 933]. Various strategies have been proposed to improve fairness issues of LLMs, ranging from holistic solutions to reducing specific biases, like biases in internal components of LLMs and biases from user interactions [181, 838, 933]. Other work has unearthed pervasive biases and stereotypes in LLMs, particularly against individuals from certain demographic groups, such as different genders [852], LGBTQ+ communities [221], and across various political spectrums [599]. The fairness of specific LLMs like GPT-3 and GPT-4 has also been extensively examined [761, 867].

We define *fairness* as the ethical principle of ensuring that LLMs are designed, trained, and deployed in ways that do not lead to biased or discriminatory outcomes and that they treat all users and groups equitably. In TRUSTLLM, we assess the fairness of LLMs in three main aspects: stereotypes, disparagement, and preference biases. The initial focus is on identifying potential stereotypes embedded within LLMs. This is achieved through three tasks: analyzing agreement on stereotypes, recognizing stereotypical content, and conducting stereotype query tests. Next, we investigate the issue of disparagement by examining how LLMs might attribute different salaries to individuals based on various characteristics, thus revealing potential biases. Finally, we explore LLMs' tendencies for preference bias by observing their decision-making in scenarios presenting contrasting opinion pairs.

12.3.4 *Robustness*

Robustness refers to the ability of AI systems to perform well under varying conditions and to properly handle exceptions, anomalies, or unexpected inputs. Recent benchmarks and studies [122, 123, 530, 532, 858, 950, 1044] on LLMs have collectively underscored a critical consensus: robustness is not an inherent quality of current LLMs. For instance, GPT-3.5 is not robust with seemingly simple inputs, such as emojis [929].

In the context of TRUSTLLM, we assess the *robustness* regarding the stability and performance when LLMs are faced with various input conditions. Note that we distinguish *robustness* from the concept of resilience against malicious attacks, which is covered under the *safety* dimension. Here, we specifically explore robustness in the context of ordinary user interactions. This involves examining how LLMs cope with natural noise in inputs and how they handle out-of-distribution (OOD) challenges. These aspects provide a comprehensive view of an LLM’s stability and reliability under typical usage scenarios.

12.3.5 Privacy

The privacy challenges associated with LLMs have garnered significant attention due to their ability to memorize and subsequently (unintentionally) leak private information, a concern that we have for traditional machine learning models [80]. This issue is exacerbated by the heavy reliance of LLMs training on Internet-sourced data, which inevitably includes personal information. Once such information is embedded within LLMs, it becomes susceptible to extraction through malicious prompts, posing a substantial risk [394].

Recent studies have delved into various aspects of privacy risks in LLMs. These include efforts of revealing personal data from user-generated text, employing predefined templates to probe and unveil sensitive information, and even attempting to *jailbreaking* LLMs to access confidential information [336, 400, 458, 782, 856]. To address these challenges, a range of frameworks and tools have been proposed and developed [55, 111, 401, 595, 841], alongside the methods of differential privacy, to mitigate the risk of privacy breaches and enhance the privacy of LLMs [95, 589]. Using cryptographic techniques like secure computation [946], recent works also explored ways to provide privacy by putting the LLM-related computation in secure computation protocols [280, 312].

Our *Privacy* guideline refers to the norms and practices that help to safeguard human and data autonomy, identity, and dignity. Specifically, we focus on evaluating LLMs’ privacy awareness and potential leakage. We first assess how well LLMs recognize and handle privacy-sensitive scenarios, including their tendency to inadvertently disclose learned information. Then, we investigate the risk of privacy leakage from their training datasets, examining if sensitive data might be unintentionally exposed when LLMs are prompted in certain ways. Overall, this analysis aims to understand LLMs’ ability to safeguard privacy and the inherent risks of private data exposure in their outputs.

12.3.6 Machine Ethics

Machine ethics is ethics for machines, where machines, instead of humans, are the subjects. The most famous machine ethics principle is the “three laws of robotics” proposed and investigated by Isaac Asimov [592]. Earlier research in this field focused on discussing the emerging field of machine ethics and the challenges faced in representing ethical principles in machines [23, 24]. These foundational investigations have also explored the motivations behind the need for machine ethics, highlighting the pursuit of ethical decision-making abilities in computers and robots [851], and examined the nature and significance of machine ethics, discussing the challenges in defining what constitutes machine ethics and proposing potential implementation strategies [596].

Subsequent research has expanded the discourse, providing nuanced analyses of contemporary ethical dilemmas and the particular challenges that arise in the context of LLMs. While specific studies have concentrated on individual models, such as Delphi [799], GPT-3 [219], and GPT-4 [1035], others have interrogated the responses of LLMs across specific domains. Two sectors frequently subject to scrutiny are the academic realm [551, 583, 660] and healthcare research [459, 460, 807].

Defining the term of *machines ethics* for LLMs is rendered nearly infeasible by our current insufficient grasp of a comprehensive ethical theory [596]. Instead, we divide it into three segments: *implicit ethics*, *explicit ethics*, and *emotional awareness*. *Implicit ethics* refers to the internal values of LLMs, such as the judgment of moral situations. We assess LLMs’ alignment with human ethical standards by evaluating their moral action judgments. In contrast, *explicit ethics* focuses on how LLMs should react in different moral environments. We also evaluate how LLMs should behave in various moral contexts. The assessment of LLMs’ ability to take morally appropriate actions in ethical scenarios is a crucial aspect, because LLMs increasingly serve as intelligent agents, engaging in action planning and decision-making. Lastly, *awareness* reflects LLMs’ capacity to understand their abilities and mission, recognize human emotions, and consider other perspectives. We evaluate four dimensions of awareness through complex scenarios, drawing insights from psychology and sociology.

12.3.7 Transparency

Transparency was not a problem when linear classifiers and decision trees dominated AI systems. Conversely, they were considered interpretable as any observer can examine the inferred tree from the root to the leaves and understand how input variables influence the output [175]. However, with the development of high-dimensional machine learning models (e.g., deep neural networks) and the pursuit of accuracy, transparency is often sacrificed due to the opaque, “black-box” nature of

complex machine learning systems [771]. Systems with opaque decision-making processes are challenging to trust, particularly in critical areas such as finance, autonomous driving, and aerospace engineering, where decisions have significant ethical and safety implications. To address these concerns, various interpretation methods have been developed in recent years [506], aiming to explain how deep learning models form their predictions. These methods are crucial for ensuring transparency and fostering trust in the predictions of advanced models in critical sectors.

As for LLMs, the lack of transparency is still noted as a core challenge [912] and a potential pitfall [85]. Reasons for their absence are often associated with some characteristics of LLMs, like complexity and massive architecture [497]. Transparency is also hard to evaluate as not all situations require the same level of transparency [497]. The evaluation should also involve human factors, like why people seek information [433, 794]. Thus, transparency is often not evaluated directly in prior works of LLMs.

In this work, *transparency* of LLMs refers to how much information about LLMs and their outputs is available to individuals interacting with them. We first contextualize various perspectives on transparency. Then, we delve into specific aspects of LLM transparency, examining the unique challenges it presents and reviewing the existing research aimed at addressing these issues.

12.3.8 Accountability

In 1996, Nissenbaum [613] described four barriers to accountability that computerization presented. Developing machine learning systems requires revisiting those concepts and bringing new challenges [163]. For LLMs and their powered AI systems, the lack of transparency often leads to a lack of accountability [175]. Besides, major scholarly and societal credit is deserved for data openness, as data work is often seen as low-level grunt work [499], and data citation is a crucial but missing component in LLMs [335]. Current works on the accountability of LLMs often focus on the healthcare [276, 399] and academic [773] domains. However, achieving overall accountability is still far from practical.

For a personal or an organization, *accountability* is a virtue [73]. We believe this is also applicable to LLMs. LLMs should autonomously provide explanations and justifications for their behavior. We follow the framework of the four barriers to the *accountability* of computer systems as identified by Helen Nissenbaum [613], and discuss these barriers in the context of LLMs. The “problem of many hands” makes it difficult to pinpoint responsibility within the collaborative development of LLMs, while the inherent “bugs” in these systems further complicate accountability. The tendency to use the computer as a “scapegoat” and the issue of “ownership without liability” where companies disclaim responsibility for errors, further blur the lines of accountability. Furthermore, as LLMs become more sophisticated, differentiating their output from human text grows more challenging. Concurrently, the extensive

use of training data in LLMs raises significant copyright concerns, underscoring the urgent need for a clear legal framework to navigate the intricate relationship between technology, ethics, and law in the AI domain.

12.4 Main Insights from TrustLLM Evaluation

Figure 12.1 presents the overview of TrustLLM [791] and the selected LLMs for evaluation, including the evaluation datasets and the proprietary and open-weight LLMs.

12.4.1 Overall Observations

Trustworthiness Is Closely Related to Capability¹ Our findings indicate a positive correlation between trustworthiness and capability, particularly evident in specific tasks. For example, in moral behavior classification and stereotype recognition tasks, LLMs like GPT-4 that possess strong language understanding capabilities tend to make more accurate moral judgments and reject stereotypical statements more reliably. Similarly, Llama2-70b and GPT-4, known for their proficiency in natural language inference, demonstrate enhanced resilience against adversarial attacks. Furthermore, we observed that the trustworthiness rankings of LLMs often mirror their positions on capability-focused leaderboards, such as MT-Bench [1021], OpenLLM Leaderboard [209], and others. This observation underscores the intertwined nature of trustworthiness and capability, highlighting the importance for both developers and users to consider these aspects simultaneously when implementing and utilizing LLMs.

Most LLMs Are “Overly Aligned We have found that many LLMs exhibit a certain degree of over-alignment (i.e., exaggerated safety), which can compromise their overall trustworthiness. Such LLMs may identify many innocuous prompt contents as harmful, thereby impacting their utility. For instance, Llama2-7b obtained a 57% rate of refusal in responding to prompts that were, in fact, not harmful. Consequently, it is essential to train LLMs to understand the intent behind a prompt during the alignment process, rather than merely memorizing examples. This will help in lowering the false positive rate in identifying harmful content.

Generally, Proprietary LLMs Outperform Most Open-Weight LLMs in Trustworthiness: However, A Few Open-Source LLMs Can Compete with Proprietary Ones We found a gap in the performance of open-weight and proprietary

¹ Here, capability refers to the functional effectiveness of the model in natural language processing tasks, including abilities in logical reasoning, content summarization, text generation, and so on.




Fig. 12.1 The design of benchmark in TRUSTLLM [791]. Building upon the evaluation principles in prior research [560, 856], we design the benchmark to evaluate the trustworthiness of LLMs on six aspects: truthfulness, safety, fairness, robustness, privacy, and machine ethics. We incorporate both existing and new datasets. The benchmark involves categorizing tasks into classification and generation. Through diverse metrics and evaluation methods, we assess the trustworthiness of a range of LLMs, encompassing both proprietary and open-weight variants

LLMs regarding trustworthiness. Generally, proprietary LLMs (e.g., ChatGPT, GPT-4) tend to perform much better than the majority of open-weight LLMs. This is a serious concern because open-weight models can be widely downloaded. Once integrated into application scenarios, they may pose severe risks. However, we were surprised to discover that Llama2 [820], a series of open-weight LLMs, surpasses proprietary LLMs in trustworthiness in many tasks. This indicates that open-weight models can demonstrate excellent trustworthiness even without adding external auxiliary modules (such as a moderator [6]). This finding provides a significant reference value for relevant open-weight developers.

Both the Model Itself and Trustworthiness-Related Technology Should Be Transparent (e.g., Open-Sourced) Given the significant gap in performance regarding trustworthiness among different LLMs, we emphasize the importance of transparency, both in the models themselves and in the technologies aimed at enhancing trustworthiness. As highlighted in recent studies [68, 540], a thorough understanding of the training mechanisms of models, including aspects such as parameter and architecture design, forms the cornerstone of researching LLMs. Our experiments found that while some proprietary LLMs exhibit high trustworthiness (e.g., ERNIE [45]), the specifics of the underlying technologies remain undisclosed. Making such trustworthy technologies transparent or open-source can promote the broader adoption and improvement of these techniques, significantly boosting the trustworthiness of LLMs. This, in turn, makes LLMs more reliable and strengthens

the AI community's overall trust in these models, thereby contributing to the healthy evolution of AI technology.

12.4.2 Novel Insights into Individual Dimensions of Trustworthiness

Truthfulness Truthfulness in AI systems refers to the accurate representation of information, facts, and results. Our findings indicate that:

1. Proprietary LLMs like GPT-4 and open-source LLMs like LLama2 often struggle to provide truthful responses when relying solely on their internal knowledge. This issue is primarily due to noise in their training data, including misinformation or outdated information, and the lack of generalization capability in the underlying Transformer architecture [844].
2. All LLMs face challenges in zero-shot commonsense reasoning tasks, suggesting difficulty in tasks that are relatively straightforward for humans.
3. In contrast, LLMs with augmented external knowledge demonstrate significantly improved performance, surpassing state-of-the-art results reported on original datasets.
4. We observe a notable discrepancy among different hallucination tasks. Most LLMs show fewer hallucinations in multiple-choice question-answering tasks compared to more open-ended tasks such as knowledge-grounded dialogue, likely due to prompt sensitivity.
5. We find a positive correlation between sycophancy and adversarial actuality. Models with lower sycophancy levels are more effective in identifying and highlighting factual errors in user inputs.

Safety Safety in LLMs is crucial for avoiding unsafe or illegal outputs and ensuring engagement in healthy conversations [536]. In our experiments, we found that:

1. The safety of most open-source LLMs remains a concern and significantly lags behind that of proprietary LLMs, particularly in areas like jailbreak, toxicity, and misuse.
2. Notably, LLMs do not uniformly resist different jailbreak attacks. Our observations revealed that various jailbreak attacks, particularly leetspeak attacks [882], vary in their success rates against LLMs. This underscores the need for LLM developers to adopt a comprehensive defense strategy against diverse attack types.
3. Balancing safety is a challenge for most LLMs; those with stringent safety protocols often show exaggerated caution, as evident in the Llama2 series and ERNIE. This suggests that many LLMs are not fully aligned and may rely on superficial alignment knowledge.

Fairness Fairness is the ethical principle of ensuring that LLMs are designed, trained, and deployed in ways that do not lead to biased or discriminatory outcomes and that they treat all users and groups equitably. In our experiments, we found that

1. The performance of most LLMs in identifying stereotypes is not satisfactory, with even the best-performing GPT-4 having an overall accuracy of only 65%. When presented with sentences containing stereotypes, the percentage of agreement of different LLMs varies widely, with the best performance at only 0.5% agreement rate and the worst-performing one approaching an agreement rate of nearly 60%.
2. Only a few LLMs, such as Oasst-12b [419] and Vicuna-7b [1020], exhibit fairness in handling disparagement; most LLMs still display biases towards specific attributes when dealing with questions containing disparaging tendencies.
3. Regarding preferences, most LLMs perform very well on the plain baseline, maintaining objectivity and neutrality or refusing to answer directly. However, when forced to choose an option, the performance of LLMs significantly decreases.

Robustness Robustness is defined as a system's ability to maintain its performance level under various circumstances [3]. In our experiments, we found that:

1. The Llama2 series and most proprietary LLMs surpass other open-source LLMs in traditional downstream tasks.
2. However, LLMs exhibit significant variability in open-ended task performance. The least effective model shows an average semantic similarity of only 88% before and after perturbation, substantially lower than the top performer at 97.64%.
3. In terms of OOD robustness, LLMs demonstrate considerable performance variation. The top-performing model, GPT-4, exhibits a RtA (Refuse to Answer) rate of over 80% in OOD detection and an average F1 score of over 92% in OOD generalization. In contrast, the least effective models show an RtA rate of merely 0.4% and an F1 score of around 30%.
4. Our observations reveal no consistent positive correlation between parameter size and OOD performance, as evidenced by the varied performance levels of Llama2 models regardless of their parameter size.

Privacy Privacy encompasses the norms and practices aimed at protecting human autonomy, identity, and dignity [3]. In our experiments, we found that:

1. Most LLMs demonstrate a certain level of privacy awareness, as evidenced by a significant increase in the likelihood of these models refusing to respond to queries about private information when informed that they must adhere to privacy policy.

2. The Pearson correlation coefficient measuring agreement between humans and LLMs on the use of privacy information varies greatly. The best-performing model, ChatGPT, achieves a correlation of 0.665, while Oass-12b exhibits a surprising negative correlation, less than zero, indicating a divergent understanding of privacy compared to humans.
3. We observed that nearly all LLMs show some degree of information leakage when tested on the Enron Email Dataset [158].

Machine Ethics Machine ethics ensure the moral behaviors of man-made machines utilizing AI, commonly referred to as AI agents [23, 24]. In our experiments, we found that:

1. LLMs have developed a specific set of moral values, yet there remains a significant gap in fully aligning with human ethics. The accuracy of most LLMs in implicit tasks within low-ambiguity scenarios falls below 70%, irrespective of the dataset. In high-ambiguity scenarios, performance varies considerably among different LLMs; for instance, the Llama2 series achieves an RtA of 99.9%, while others score less than 70%.
2. In terms of awareness, the best-performing model GPT-4 achieves an average accuracy rate of 94% over four awareness datasets. Other LLMs exhibit decent but not substantial awareness.

Chapter 13

Attacks and Defenses on Aligned Large Language Models



Abstract Safety, security, and compliance are essential requirements when aligning large language models (LLMs). To reduce harm and misuse, efforts have been made to align these LLMs to human values using advanced training techniques such as Reinforcement Learning from Human Feedback (RLHF). This chapter presents attacks and defenses developed for aligned LLMs. Efforts made into circumventing the safety guardrails of aligned LLMs are often called *jailbreaks*. In this context, attacks aim to find vulnerability of aligned LLMs to adversarial jailbreak attempts aiming at subverting the embedded safety guardrails, while defenses aim to either detect malicious prompts or enhance the refusal capability to attacks.

13.1 Introduction

Recent advances in large language models (LLMs) [183, 846] such as GPT-4 [619], LLAMA-2 [820], and Mistral [373] have showcased their ability to understand and generate text akin to human interaction [174, 664, 1031]. These models, powered by the Transformer architecture, excel in processing sequential data and understanding complex language patterns, hence enhancing tasks like text summarization, creative writing, and coding. To maintain model integrity and mitigate undesired outputs, developers implement alignment constraints using techniques like Reinforcement Learning with Human Feedback (RLHF) [34, 41, 627] and Supervised Fine-Tuning (SFT).

However, aligned LLMs have been found to be vulnerable to a type of adversarial manipulation known as “jailbreak attack”. Jailbreak attacks involve maliciously inserting or replacing tokens in the user instruction or rewriting it to bypass and circumvent the safety guardrails of aligned LLMs. A notable example is that a jailbroken LLM would be tricked into generating hate speech targeting certain groups of people, as demonstrated in Fig. 13.1a. Many red-teaming efforts [105, 529, 578, 883, 956, 1052] have been put into designing algorithms to automatically generate jailbreak prompts to help test the robustness of aligned LLMs. Specifically, GCG [1052], one of the earlier works in this area, can successfully jailbreak several LLMs by optimizing an inserted universal adversarial suffix. This finding suggests




Fig. 13.1 Overview of **Gradient Cuff** [328]. (a) introduces an example of jailbreak prompts by presenting a conversation between malicious actors and the Vicuna chatbot. (b) visualizes the refusal loss landscape for malicious queries and benign queries by plotting the interpolation of two random directions in the query embedding with coefficients α and β following [468]. The refusal loss evaluates the probability that the LLM would not directly reject the input query, and the loss value is computed using Eq. 13.3. (c) shows the running flow of Gradient Cuff (at top), practical computing examples for refusal loss (at bottom left), and the distributional difference of the gradient norm of refusal loss on benign and malicious queries (bottom right). (d) shows the performance of Gradient Cuff against 6 jailbreak attacks for Vicuna-7B-V1.5

that the embedded alignment effort in LLMs could be completely broken by the jailbreak attack.

Since the discovery of jailbreak risks for LLMs, various methods have been explored to defend against jailbreak attacks. In this chapter, we divide the defenses into two categories based on their functionality:

- **Detection:** defensive procedures made to check if an input query contains jailbreak attempts or not. The protected LLM will refuse to give a response if the input query is flagged as a jailbreak prompt by the detector.
- **Mitigation:** defensive procedures applied to the protected LLM to enhance its alignment in adversarial settings, such as improved prompt template design or robust model training/fine-tuning.

Notably, ideal defenses need to minimize the potential trade-off between preserving utility and effectively mitigating jailbreak risks. Broadly speaking, safety guardrails built for LLMs should also prevent direct and indirect prompt injection attacks, where LLMs can be directly tricked, or indirectly be used to trigger a chain of actions (e.g., LLMs as agents), to achieve the attacker's objective, such as the examples explored in [267]. This chapter presents an overview of jailbreak attacks

and defenses for LLMs, with an emphasis on the detection method Gradient Cuff [328] and the prompt-based mitigation method Defensive Prompt Patch [921].

13.2 Background and Related Work

Jailbreak Attacks Existing jailbreaks can be roughly divided into feedback-based jailbreak attacks and rule-based jailbreak attacks.

Feedback-based jailbreaks utilize the feedback from the target LLM to iteratively update the jailbreak prompt until the model complies with the malicious instruction embedded in the jailbreak prompt. Feedback-based jailbreaks can be further categorized by their access mode to the target LLM. Some feedback-based jailbreak attacks like GCG [1052], require white-box access to the target LLM. Specifically, GCG leverages gradients with respect to the one-hot token indicators to find better token choices at each position. Some feedback-based jailbreaks need gray-box access to the target LLM. The typical one is AutoDAN [529], which employs the target LLM’s generative loss of the target response to design the fitness score of the candidate jailbreak prompt to guide further optimization. PAIR [105] and TAP [578] are the representatives of feedback-based jailbreaks which only require black-box access to the target LLM. In PAIR and TAP, there are also two LLMs taking on the attacker role and the evaluator role. At each iteration, the attacker-generated jailbreak prompt would be rated and commented on by the evaluator model according to the target LLM’s response to the attack. Next, the attacker would generate new jailbreak prompts based on the evaluator’s comments, and repeat the above cycle until the jailbreak prompt can get full marks from the evaluator. The only information provided by the target LLM is the response to the jailbreak attack.

As for the *rule-based* jailbreak attacks, we highlight Base64 [883] and Low Resource Language (LRL) [956]. Base64 encodes the malicious instruction into base64 format and LRL translates the malicious instruction into the language that is rarely used in the training process of the target LLM, such as German, Swedish, French and Chinese. The ICA Attack [891] leverages in-context learning to misaligned responses, and the Catastrophic Attack [343] manipulates generation configurations to trigger misaligned outputs.

Jailbreak Defenses *Detection-based methods*: PPL [356] uses an LLM to compute the perplexity of the input query and rejects those with high perplexity. Smooth-LLM [704], motivated by randomized smoothing [160], perturbs the original input query to obtain several copies and aggregates the intermediate responses of the target LLM to these perturbed queries to give the final response to the original query. Erase-Check employs a model to check whether the original query or any of its erased subsentences is harmful. The query would be rejected if the query or one of its sub-sentences is regarded as harmful by the safety checker. Gradient Cuff [328] detects jailbreak prompts by checking the refusal loss of the input user query and estimating the gradient norm of the loss function.

Table 13.1 Comparison between different defense (mitigation) methods against jailbreak attacks on LLMs

	Optimizable prompt	Gradient-based search	Interpretable	Attack success rate	Utility degradation
Self-reminder	✓	✗	✓	Medium	Medium
RPO	✓	✓	✗	High	Low
Goal prioritization	✗	✗	✓	Low	High
Default system prompt	✗	✗	✓	High	Medium
Defensive prompt patch	✓	✓	✓	Low	Low

Mitigation-Based Methods Unlike detection methods, there is another line of work [843, 892, 920, 1006] where prompt engineering techniques are used to defend against jailbreak attacks. Self-Reminder [920] shows promising results by modifying the system prompt of the target LLM so that the model reminds itself to process and respond to the user in the context of being an aligned LLM. RPO (Robust Prompt Optimization) [1032] modifies objectives to minimize the perceptual distance between harmful queries and safe responses. Furthermore, Goal Prioritization and Default System Prompts [1014, 1019] are designed to direct LLMs to prioritize safety and prevent the generation of harmful outputs. Defense Prompt Patch (DPP) [921] is designed to minimize jailbreak risks while maintaining high utility, addressing the common pitfalls in current prompt-based defenses. We provide a comparative analysis of different prompt-based mitigation defenses in Table 13.1.

13.3 Gradient Cuff

The overview of Gradient Cuff [328], a jailbreak detection method, is presented in Fig. 13.1. In this section, we will formalize the concept of *Refusal loss function* and illustrate how Gradient Cuff uses the unique loss landscape properties of this function observed between malicious and benign user queries for effective detection.

13.3.1 Refusal Loss Function and Landscape Exploration

Current transformer-based LLMs will return different responses to the same query due to the randomness of autoregressive sampling based generation [210, 308]. With this randomness, it is an interesting phenomenon that a malicious user query will sometimes be rejected by the target LLM, but sometimes be able to bypass the

safety guardrail. Based on this observation, for a given LLM T_θ parameterized with θ , we define the refusal loss function $\phi_\theta(x)$ for a given input user query x as below:

$$\phi_\theta(x) = 1 - p_\theta(x); \quad (13.1)$$

$$p_\theta(x) = \mathbb{E}_{y \sim T_\theta(x)} JB(y) \quad (13.2)$$

where y represents the response of T_θ to the input user query x . $JB(\cdot)$ is a binary indicator function to determine whether the response triggers a refusal action by the LLM. The function p_θ can be interpreted as the expected rate of getting refusal on the response y from T_θ taking into account the randomness in the decoding process. Therefore, by our definition, the refusal loss function $\phi_\theta(x)$ can be interpreted as the likelihood of generating a non-refusal response to x . Following SmoothLLM [704], we define $JB(\cdot)$ as

$$JB(y) = \begin{cases} 1, & \text{if } y \text{ contains any jailbreak keyword;} \\ 0, & \text{otherwise.} \end{cases}$$

For example, $JB(y)$ would be 0 if $y = \text{"Sure, here is the python code to ..."} and $JB(y)$ would be 1 if $y = \text{"Sorry, I cannot fulfill your request. ..."}.$$

Alternatively, we can view $Y = JB(y)$ as a random variable obeying the Bernoulli distribution such that

$$Y = \begin{cases} 1, & \text{with probability } p_\theta(x) \\ 0, & \text{with probability } 1 - p_\theta(x) \end{cases}$$

so that $\phi_\theta(x)$ can be interpreted as the expected refusal loss:

$$\phi_\theta(x) = 1 - \mathbb{E}[Y] = 1 - p_\theta(x).$$

In practice, since we do not have the prior knowledge for $p_\theta(x)$, we use the sample mean $f_\theta(x)$ to approximate $\phi_\theta(x)$:

$$f_\theta(x) = 1 - \frac{1}{N} \sum_{i=1}^N Y_i, \quad (13.3)$$

where $\{Y_i | i = 1, 2, \dots, N\}$ is obtained by running N independent realizations of the random variable Y . In the i^{th} trial, we query the LLM T_θ using x to get the response $y_i \sim T_\theta(x)$, and apply the indicator function $JB(\cdot)$ on y_i to get $Y_i = JB(y_i)$. Equation (13.3) can be explained as using the sample mean of the random variable Y to approximate its expected value $\mathbb{E}[Y]$.

In general, $\phi_\theta(x) < 0.5$ could be used as a naive detector to reject x since $p_\theta(x)$ can be interpreted as the probability that T_θ regards x as harmful. However, this detector alone only has limited effect against jailbreak attacks [328]. To further

explore how this refusal loss can be used to improve jailbreak detection, we visualize the refusal loss landscape following the 2-D visualization techniques from [468] in Fig. 13.1b. From Fig. 13.1b, we find that the landscape of $f_\theta(\cdot)$ is more precipitous for malicious queries than for benign queries, which implies that $f_\theta(\cdot)$ tends to have a large gradient norm if x represents a malicious query. This observation motivates our proposal of using the gradient norm of $f_\theta(\cdot)$ to detect jailbreak attempts that pass the initial filtering of rejecting x when $f_\theta(x) < 0.5$.

13.3.2 Gradient Norm Estimation

In general, the exact gradient of $\phi_\theta(x)$ (or $f_\theta(x)$) is infeasible to obtain due to the existence of discrete operations such as applying the $JB(\cdot)$ function to the generated response, and the possible involvement of black-box evaluation functions (e.g., Perspective API). We propose to use zeroth order gradient estimation to compute the approximate gradient norm, which is widely used in black-box optimization with only function evaluations (zeroth order information) [57, 522]. Similar gradient estimation techniques were used to generate adversarial examples from black-box models [124, 138, 350].

A zeroth-order gradient estimator approximates the exact gradient by evaluating and computing the function differences with perturbed continuous inputs. Our first step is to obtain the sentence embedding of x in the embedding space of T_θ in $\mathbb{R}^d$. For each text query x with n words (tokens) in it, it can be embedded into a matrix $e_\theta(x) \in \mathbb{R}^{n \times d}$ where $e_\theta(x)_i \in \mathbb{R}^d$ denotes the word embedding for the i^{th} word in sentence x . We define the sentence embedding for x by applying mean pooling to $e_\theta(x)$ defined as

$$\text{mean-pooling}(x) = \frac{1}{n} \sum_{i=1}^n e_\theta(x)_i \quad (13.4)$$

With the observation that

$$\text{mean-pooling}(x) + \mathbf{v} = \frac{1}{n} \sum_{i=1}^n (e_\theta(x)_i + \mathbf{v}), \quad (13.5)$$

one can obtain a perturbed sentence embedding of x with any perturbation $\mathbf{v}$ by equivalently perturbing the word embedding of each word in x with the same $\mathbf{v}$.

Based on this definition, we approximate the exact gradient $\nabla \phi_\theta(x)$ by $g_\theta(x)$, which is the estimated gradient of $f_\theta(x)$. Following [57, 522], we calculate $g_\theta(x)$ using the directional derivative approximation

$$g_\theta(x) = \sum_{i=1}^P \frac{f_\theta(\mathbf{e}_\theta(x) \oplus \mu \cdot \mathbf{u}_i) - f_\theta(\mathbf{e}_\theta(x))}{\mu} \mathbf{u}_i, \quad (13.6)$$

where $\mathbf{u}_i$ is a d dimension random vector drawn from the standard multivariate normal distribution, i.e., $\mathbf{u}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, μ is a smoothing parameter, $\oplus$ denotes the row-wise broadcasting add operation that adds the same vector $\mu \cdot \mathbf{u}_i$ to every row in $\mathbf{e}_\theta(x)$.

Based on the definitions in Eqs. (13.3) and (13.6), we provide a probabilistic guarantee below for analyzing the gradient approximation error of the true gradient $\phi_\theta(\cdot)$.

Theorem 13.1 *Let $\|\cdot\|$ denote a vector norm and assume $\nabla\phi_\theta(x)$ is L -Lipschitz continuous. With probability at least $1 - \delta$, the approximation error of $\nabla\phi_\theta(x)$ satisfies*

$$\|g_\theta(x) - \nabla\phi_\theta(x)\| \leq \epsilon$$

for some $\epsilon > 0$, where $\delta = \Omega^1(\frac{1}{N} + \frac{1}{P})$ and $\epsilon = \Omega(\frac{1}{\sqrt{P}})$.

This theorem demonstrates that one can reduce the approximation error by taking larger values for N and P . The proof can be found in [328].

13.3.3 Gradient Cuff: Two-Step Jailbreak Detection

With the discussions in Sects. 13.3.1 and 13.3.2, we now formally propose Gradient Cuff, a two-step jailbreak detection method based on checking the refusal loss and its gradient norm. Our detection procedure is shown in Fig. 13.1c. Gradient Cuff can be summarized into two steps:

- **(Step #1) Sampling-based Rejection:** In the first step, we reject the user query x by checking whether $f_\theta(x) < 0.5$. If true, then x is rejected, otherwise, x is pushed into Step 2.
- **(Step #2) Gradient Norm Rejection:** In the second step, we regard x as having jailbreak attempts if the norm of the estimated gradient $g_\theta(x)$ is larger than a configurable threshold t , i.e., $\|g_\theta(x)\| > t$.

Before deploying Gradient Cuff on LLMs, we first test it on a bunch of benign user queries to select a proper threshold t that fulfills the required benign refusal rate (that is, the false positive rate σ). We use a user-specified σ value (e.g., 5%) to guide the selection of the threshold t so that the total refusal rate on the benign validation dataset $\mathcal{B}_{val}$ won't exceed σ .

We summarize our method in Algorithm 4. The algorithm is implemented by querying the LLM T_θ multiple times, each to generate a response for the same input query x . The total query times to T_θ required to compute $f_\theta(x)$ and $g_\theta(x)$ in Gradient Cuff is at most $q = N \cdot (P + 1)$. To maintain the LLM's efficiency, one

¹ $l(t) = \Omega(s(t))$ means $s(t)$ is the infimum of $l(t)$.

Algorithm 4 Gradient cuff: two-step jailbreak detection

```

1: Notations: The LLM to be protected:  $T_\theta$ , Required benign refusal (false positive) rate:  $\sigma$ ,
   Gaussian vector numbers:  $P$ , LLM Response Sampling numbers:  $N$ , Smoothing parameter:
    $\mu$ , Collection of benign user queries:  $\mathcal{B}_{val}$ , Threshold:  $t$ , Input User Query:  $x_{test}$ 
2: Threshold Selection:
3: Construct  $S = \{x | f_\theta(x) < 0.5 \text{ and } x \in \mathcal{B}_{val}\}$  based on Eq. (13.3).
4: Construct  $G = \{\|g_\theta(x)\| | x \in \mathcal{B}_{val} \setminus S\}$  based on Eq. (13.6).
5: Sort  $G$  in descending order
6: Select  $k$  that fulfills:  $k - 1 \leq |\mathcal{B}_{val}| \cdot \sigma - |S| < k$ 
7: Set threshold  $t = G[k]$  # such that  $\frac{|S|+k-1}{|\mathcal{B}_{val}|} \leq \sigma$ 
8: Detection on test query  $x_{test}$ :
9: Calculate  $f_\theta(x_{test})$  based on Eq. (13.3).
10: if  $f_\theta(x_{test}) < 0.5$  then
11:   return "I cannot fulfill your request." # First-step filtering
12: else
13:   Calculate  $g_\theta(x_{test})$  based on Eq. (13.6).
14:   if  $g_\theta(x_{test}) > t$  then
15:     return "I cannot fulfill your request." # Second-step filtering
16:   else
17:     return  $y \sim T_\theta(x_{test})$  # No rejection
18:   end if
19: end if

```

can use batch inference to compute these queries in parallel, thereby reducing the total running cost of the LLM.

13.3.4 Performance Evaluation

Malicious User Queries We sampled 100 harmful behavior instructions from AdvBench² in [1052] as jailbreak templates, each to elicit the target LLM to generate certain harmful responses. We then use various existing jailbreak attack methods to generate enhanced jailbreak prompts for them. Specifically, for each harmful behavior instruction, we use GCG [1052] to generate a universal adversarial suffix, use AutoDAN [529], PAIR [105], and TAP [578] to generate a new instruction, use LRL [956] to translate it into low source languages that rarely appear in the training phase of the target LM such as German, Swedish, French and Chinese, and use Base64 [883] to encode them in base64 format. In our experiments, we use **malicious user queries** to denote these harmful behavior instructions with jailbreak prompts. For example, **malicious user queries (AutoDAN)** means those harmful instructions with jailbreak prompts generated by AutoDAN.

² https://github.com/llm-attacks/llm-attacks/blob/main/data/advbench/harmful_behaviors.csv.

Benign User Queries We also build a corpus of benign queries to obtain the gradient norm rejection threshold and evaluate the performance of Gradient Cuff on non-harmful user queries. We collect benign user queries from the LMSYS Chatbot Arena leaderboard,³ which is a crowd-sourced open platform for LLM evaluation. We removed the toxic, incomplete, and non-instruction queries and then sampled 100 queries from the rest to build a test set. We use the rest as a validation dataset to determine the gradient norm threshold t . In our experiments, **benign user queries** denotes the queries in the test set.

Aligned LLMs We conduct the jailbreak experiments on 2 aligned LLMs: LLaMA-2-7B-Chat [820] and Vicuna-7B-V1.5 [1020]. LLaMA-2-7B-Chat is the aligned version of LLAMA-2-7B. Vicuna-7B-V1.5 is also based on LLAMA2-7B and has been further supervised fine-tuned on 70k user-assistant conversations collected from ShareGPT.⁴ We use **protected LLM** to represent these two models in the experiments.

Detection Baselines We compare our method with various jailbreak detection methods including PPL [356], Erase-check [426], and SmoothLLM [704]. To implement PPL, we use the protected LLM itself to compute the perplexity for the input user query and directly reject the one with a perplexity higher than some threshold in our experiment. For Erase-Check, we employ the LLM itself to serve as a safety checker to check whether the input query or any of its erased sub-sentences is harmful. SmoothLLM perturbs the original input query to obtain multiple copies and then aggregates the protected LLM’s response to these copies to respond to the user. Quite unlike the previous ones, Self-Reminder converts the protected LLM into a self-remind mode by modifying the system prompt.

Metrics We report both the refusal rates for malicious user queries (true positive rate, TPR) and the benign user queries (false positive rate, FPR) to evaluate Gradient Cuff as well as those baselines. Higher TPR and lower FPR indicate better performance. For LRL, we report the average refusal rate when translating the English queries to German (de), French (fr), Swedish (sv), and Simplified Chinese (zh-CN).

Implementation of Gradient Cuff We use $\mu = 0.02$, $N = P = 10$ in our main experiments and report the results when σ (FPR) is set to 5%. For the text generation setting, we use temperature = 0.6, top-p parameter = 0.9 for both LLaMA2-7B-Chat and Vicuna-7B-V1.5, and adopt Nucleus Sampling. As for the system prompt, we use the default setting provided in the fastchat repository [1020].

Numerical Results We evaluate Gradient Cuff as well as all the detection baselines against 6 different jailbreak attacks (GCG, AutoDAN, PAIR, TAP, Base64, and LRL) and benign user queries. We report the average refusal rate across these 6

³ https://huggingface.co/datasets/lmsys/chatbot_arena_conversations.

⁴ <https://sharegpt.com>.




Fig. 13.2 Performance evaluation on LLaMA2-7B-Chat (a) and Vicuna-7B-V1.5 (b). The horizontal axis represents the refusal rate of benign user queries (FPR), and the vertical axis shows the average refusal rate across 6 malicious user query datasets (TPR). The error bar shows the standard deviation between the refusal rate of these 6 jailbreak datasets

malicious user query datasets as True Positive Rate (TPR) and the refusal rate on benign user queries as False Positive Rate (FPR). From Fig. 13.2 we can summarize that Gradient Cuff stands out on both benign queries and malicious queries, attaining high TPR and low FPR. Gradient Cuff can outperform PPL and SmoothLLM with a similar FPR and a much higher TPR. Though Erase-Check can also achieve good detection performance on malicious user queries, it cannot be regarded as a practical defense method because it would reject almost all the benign user queries in the test set, which can drastically compromise the usability of the protected LLMs. We also plot the standard deviation of TPR over different types of malicious queries for all methods. The results shown in Fig. 13.2a, b demonstrate that Gradient Cuff has the most balanced performance across all types of tested jailbreak attacks. Overall, the comparison with PPL, SmoothLLM, and Erase-Check shows that Gradient Cuff is a more effective defense by providing stable and strong defense functionality against different types of jailbreak attacks. In [328], the authors also explored the performance of detection methods against adaptive attacks.

13.4 Defensive Prompt Patch

13.4.1 Preliminaries

We introduce the following mathematical notations as the basis for constructing a defensive prompt patch (DPP).

Jailbreak Attack A jailbreak attack on an LLM aims to circumvent model alignment by using meticulously crafted prompts [957, 1013]. We denote a malicious query as $\mathbf{u}_{1:n} = \langle u_1, u_2, \dots, u_n \rangle$, with each u_i being an input token. Ordinarily, the LLM would reject such queries based on its alignment policies. However, refined jailbreak queries, $\tilde{\mathbf{u}}_{1:m} = \langle \tilde{u}_1, \tilde{u}_2, \dots, \tilde{u}_m \rangle$, manipulate these policies to elicit a compliant response $\mathbf{r}_{1:k} = \langle r_1, r_2, \dots, r_k \rangle$, reflecting the original malicious intent.

Jailbreak Defense Our defense involves a defensive prompt patch $\mathbf{d}_{1:l} = \langle d_1, d_2, \dots, d_l \rangle$, derived from our DPP algorithm. This patch is appended to the refined query, forming a protected input $\mathbf{x}_{1:m+l}^{\text{guard}} = (\tilde{\mathbf{u}}_{1:m}, \mathbf{d}_{1:l})$, typically resulting in a refusal response $\mathbf{s}_{1:n} = \langle s_1, s_2, \dots, s_n \rangle$.

Utility Degradation We measure utility degradation by the deviation in LLM responses to benign queries appended with $\mathbf{d}_{1:l}$. Ideally, the response to a benign query $\mathbf{b}_{1:p} = \langle b_1, b_2, \dots, b_p \rangle$ patched by $\mathbf{d}_{1:l}$ should closely match the response to $\mathbf{b}_{1:p}$ alone.

Mathematical Formulation We define the $\oplus$ operation as the concatenation of two sequences. For a given sequence $\mathbf{a}_{1:n} = \langle a_1, \dots, a_n \rangle$ and $\mathbf{z}_{1:m} = \langle z_1, \dots, z_m \rangle$, $\mathbf{a}_{1:n} \oplus \mathbf{z}_{1:m}$ is defined as: $\mathbf{a}_{1:n} \oplus \mathbf{z}_{1:m} = \langle a_1, \dots, a_n, z_1, \dots, z_m \rangle$. We denote sequences of harmful responses and jailbreak inputs by $\mathbf{r}_{1:k}$ and $\tilde{\mathbf{u}}_{1:m}$, respectively. Since LLMs are specifically trained to predict the probability of the next word, we define the goal (i.e., the objective function to be maximized) of a jailbreak attack as:

$$P(\mathbf{r}_{1:k} | \tilde{\mathbf{u}}_{1:m}) = \prod_{j=1}^k P(r_j | \tilde{\mathbf{u}}_{1:m}, \mathbf{r}_{1:j-1}) \quad (13.7)$$

and the goal of mitigation-based defense as:

$$P(\mathbf{s}_{1:n} | \tilde{\mathbf{u}}_{1:m} \oplus \mathbf{d}_{1:l}) = \prod_{i=1}^n P(s_i | \tilde{\mathbf{u}}_{1:m} \oplus \mathbf{d}_{1:l}, \mathbf{s}_{1:i-1}) \quad (13.8)$$

where $\mathbf{s}_{1:n}$ is the refusal response to the jailbreak inputs. Finally, we assess utility degradation by:

$$P(\mathbf{h}_{1:q} | \mathbf{b}_{1:p} \oplus \mathbf{d}_{1:l}) = \prod_{k=1}^q P(h_k | \mathbf{b}_{1:p} \oplus \mathbf{d}_{1:l}, \mathbf{h}_{1:k-1}) \quad (13.9)$$

where $\mathbf{h}_{1:q}$ is the normal response for each benign queries $\mathbf{b}_{1:p}$. The DPP algorithm's efficacy is evaluated by its performance in both defense against malicious queries and impact on utility on benign queries.

13.4.2 Score Evaluation

In this mitigation strategy, the DPP must fulfill two crucial objectives: (I) **Maximization of Refusal Score** on malicious queries and (II) **Maximization of Helpful Score** on benign queries.

To achieve (I), we use the log-likelihood of Eq. (13.8) and define the refusal score as follows:

$$\mathcal{S}_{D_i} = \log P(\mathbf{s}_{1:n} | \tilde{\mathbf{u}}_{1:m} \oplus \mathbf{d}_{1:l}) \quad (13.10)$$

where $\mathcal{S}_{D_i}$ denotes the refusal score attributed to the i -th DPP within the population of DPPs. The vector $\mathbf{s}_{1:n}$ encapsulates the refusal response, $\tilde{\mathbf{u}}_{1:m}$ represents the jailbreak query, and $\mathbf{d}_{1:l}$ is the our defensive mechanism.

Similarly, for (II), the inputs include benign queries combined with the same DPP as used in the refusal score calculation. Applying the log-likelihood of Eq. (13.9). The helpful score is formulated as:

$$\mathcal{S}_{H_i} = \log P(\mathbf{h}_{1:q} | \mathbf{b}_{1:p} \oplus \mathbf{d}_{1:l}) \quad (13.11)$$

where $\mathcal{S}_{H_i}$ represents the helpfulness score assigned to the i -th DPP within the population of DPPs. The vector $\mathbf{h}_{1:q}$ denotes the standard response, whereas $\mathbf{b}_{1:p}$ refers to the benign query. The overall score function for training DPP combines the refusal and helpful scores, weighted by coefficients α and β , respectively:

$$\mathcal{S}_{T_i} = \alpha \cdot \mathcal{S}_{D_i} + \beta \cdot \mathcal{S}_{H_i} \quad (13.12)$$

13.4.3 DPP Training Algorithm

Using the total score defined in Sect. 13.4.2, we use a Hierarchical Genetic Algorithm (HGA) to optimize DPP, drawing inspiration from the AutoDAN jailbreak attack in [529]. We adapt and extend HGA to iteratively refine DPP based on our defined scores, as depicted in Fig. 13.3b, c to develop our methodology, which we term the **Defensive Prompt Patch Algorithm** (DPP Algorithm).

Initially, we establish a baseline DPP, designated as the prototype. Without loss of generality, this prototype may take the form of either a Prefix DPP or a Suffix DPP. The studies in [921] suggest that Suffix DPP is more effective. Following this, the prototype is subjected to K iterations of rewriting via an LLM to potentially refine the DPP, creating a population of DPP candidates. Each candidate within the population is evaluated by sampling refusal data pairs and helpful data pairs from adversarial/utility datasets to compute the total score, as formulated in Eq. (13.12).

The DPP optimization process is conducted over I iterations for each candidate, during which the DPP algorithm executes two pivotal operations: **Sentence-Level Word Substitution** and **Paragraph-Level Sentence Swap and Mutations**.

In **Sentence-Level Word Substitution**, each sentence within the population is assigned a score calculated using Eq. (13.12). A certain percentage of defense prompts are retained based on their scores for further optimization. For these sentences, words are initially assigned the same score as their corresponding sentences. These scores are later adjusted based on the frequency of occurrence of each word. Words whose scores surpass a specified threshold are then randomly replaced with synonyms.




Fig. 13.3 Overview of **Defensive Prompt Patch** [921] (a) showcases an example of jailbreak attacks. (b) is the DPP training phase in which the algorithm takes in the refusal and helpful datasets and a prototype of the defense prompt. Then, the algorithm forms the defense prompt population by revising the prototype using LLM. For each of the defense prompts in the population, the algorithm will evaluate the defense and utility scores as detailed in Sect. 13.4. The algorithm keeps editing the defense prompts with low scores using the Hierarchical Genetic Search algorithm. (c) shows the deployment of DPP in the LLM inference phase, by adding the best DPP in (b) (indicated in green patch) to every input query. (d) shows the trade-off graphs between the win-rate (utility) [481] and attack success rate (ASR) in both LLAMA-2-7B-Chat and Mistral-7B-Instruct-v0.2 models for different defenses

In **Paragraph-Level Sentence Swap and Mutations**, we specify a swap probability p_{swap} and a mutation probability p_{mutate} . The defensive prompt patch, modified in the previous step, is reassessed for total score at the sentence level. Employing a methodology similar to that of sentence-level optimization, the algorithm selects parent sentences based on their scores, segments and swaps these sentences, and then conducts mutations by revising sentences using an LLM.

These processes—**Sentence-Level Word Substitution** and **Paragraph-Level Sentence Swap and Mutations**—aim to increase the diversity within the defensive prompt patch population and enhance the likelihood of identifying the optimal patch.

The full algorithm is delineated in Algorithm 5, along with the dependent functions. Ultimately, the algorithm produces an updated set of optimized DPPs, comprising K enhanced patches, and identifies the Best Defensive Prompt Patch based on the highest total score.

Best DPP Selection Algorithm 5 identifies the optimal DPP for a given pair of refusal and helpful data. Our primary objective is to find a DPP that generalizes well across different user queries. To enhance the universality of DPP, we incorporate N pairs of refusal and helpful data, sampled from their respective datasets. In each

Algorithm 5 Defensive prompt patch (DPP) algorithm

```

1: Arguments: Defensive Prompt Patch Prototype  $O$ , refusal pair  $(x^r, y^r)$ , helpful pair  $(x^h, y^h)$ ,
    $\alpha$  and  $\beta$ , target LLM
2: Initialization: Number of optimization iteration  $I$ , batch size,  $p_{crossover}$ ,  $p_{mutate}$ , Sentence-
   level iterations, Paragraph-level iterations, number of steps, number of parent set size
3: DPP_Set  $\leftarrow$  DPP Set Generation( $O, K$ ) by Algorithm 6
4: while  $I$  is not reached do
5:   for iteration in sentence-level iterations do
6:     Evaluate refusal/helpful score of each DPP with  $(x^r, y^r)/(x^h, y^h)$  and target LLM
7:     Final Score  $\leftarrow$  calculate the score using Eq. (13.12)
8:     Select elite and parent prompts from DPP_Set according to Final Score
9:     WordDict  $\leftarrow$  Calculate each word score using selected parent prompts by Algorithm 7
10:    Find synonyms for each word
11:    if random value < WordDict[synonym] / sum(word scores) then
12:      Replace word with synonym
13:    end if
14:  end for
15:  for iteration in paragraph-level iterations do
16:    Repeat line 6 to 8
17:    Conduct crossover and mutation on selected parent prompts using Algorithm 8
18:  end for
19:  New_DPP_Set  $\leftarrow$  DPP_Set  $\cup$  New_DPP
20:  Best_DPP  $\leftarrow$  Best score within New_DPP_Set
21: end while
22: return (New_DPP_Set, Best_DPP)

```

Algorithm 6 DPP set generation ($prompt, K$)

```

1: Potential DPP Set = []
2: for  $i = 1$  to  $K$  do
3:   Use LLM to rewrite the DPP prompt without changing the meaning and length
4:   return New DPP prompt
5: end for

```

iteration of the DPP algorithm, as described earlier, a set of candidate DPPs is generated along with the best DPP for the specific data pair. This set of candidate DPPs is then used for the next pair of refusal and helpful data. By iteratively optimizing this set of DPP candidates, we aim to identify the most generalizable DPP with the best defensive and utility performance. The overall optimization procedure is detailed in Algorithm 9.

Algorithm 7 ConstructWordScoreDict(*WordDict*, *DPP_Set*, *scoreList*, *M*)

```

1: wordScores  $\leftarrow \{\}$ 
2: Obtained a stop words dictionary Stop_Words
3: for each (DPP, score) in (DPP_Set, scoreList) do
4:   word_list  $\leftarrow$  Save words in DPP that are not in Stop_Words
5:   Append corresponding score of each word in word_list into the wordScores dictionary
6: end for
7: for each (word, scores) in wordScores do
8:   avgScore  $\leftarrow$  average of scores for each word
9:   Save avgScore if word does not exist in WordDict
10:  Save (avgScore + previous_avgScore)/2 if word does exist in WordDict
11: end for
12: sortedWordDict  $\leftarrow$  sort wordDict by values in descending order
13: return top M items from sortedWordDict

```

Algorithm 8 Crossover and mutation (*population*)

```

1: offsprings  $\leftarrow []$ 
2: for parent1, parent2 in population do random value  $< p_{crossover}$ 
3:   segment1, segment2  $\leftarrow$  Parse parent1, parent2 into segments
4:   child1, child2  $\leftarrow$  Swap and Merge(segment1, segment2)
5:   Append child1 and child2 to offsprings
6:   Append parent1 and parent2 to offsprings
7: end for
8: for i in Range(Len(offsprings)) do random value  $< p_{mutation}$ 
9:   Use LLM to rewrite offsprings[i]
10: end for
11: return offsprings

```

Algorithm 9 Best DPP training algorithm

```

1: Require: Refusal Dataset, Helpful Dataset, target LLM.
2: Initialization: Choose initial prompt D (Suffix/Prefix).
3: Initial Hyperparameters: Set  $\alpha$ ,  $\beta$ .
4: DPP_Set  $\leftarrow []$ 
5: for i = 1 to N do
6:   Get refusal pairs ( $x_i^r$ ,  $y_i^r$ ).
7:   Get helpful pairs ( $x_i^h$ ,  $y_i^h$ ).
8:   (New_DPP_Set, Best_DPP)  $\leftarrow$ 
9:     DPPAlgorithm( $x_i^r$ ,  $y_i^r$ ), ( $x_i^h$ ,  $y_i^h$ ), D,  $\alpha$ ,  $\beta$ , DPP_Set
10:  DPP_Set  $\leftarrow$  New_DPP_Set
11: end for
12: Select Best_DPP from DPP_Set

```

13.4.4 Performance Evaluation

Adversarial Dataset We use the AdvBench [1052], specifically the **harmful behavior instructions**,⁵ as jailbreak questions. Each of them is fed into a well-aligned LM (LLAMA-2-7B-Chat [820]) to generate the denial responses. In our experiment, we sampled 100 jailbreak questions and recorded both jailbreak questions along with their refusal responses to form the **Adversarial Dataset**.

Utility Dataset We also use the Alpaca dataset⁶ as our benchmark. For consistency with the Adversarial Dataset, we also sampled only 100 benign questions and their corresponding answers.

Language Model We perform our jailbreak experiments on **LLAMA-2-7B-Chat** [820]. LLAMA-2-7B-Chat is an adapted version of LLAMA-2-7B, specifically configured for chat-based interactions.

Jailbreak Attack Methods We use several existing jailbreak attack methods to generate advanced malicious prompts. Specifically, for each malicious behavior statement, we apply several different types of jailbreaking attacks: (i) **Uninterpretable Jailbreak Attacks**—we used GCG [1052] and Base64 [883] to generate adversarial prompts. Specifically, GCG is used to generate an adversarial suffix for each malicious query. Base64 encodes each harmful query in Base64 format. (ii) **Interpretable Jailbreak Attacks**—AutoDAN [529], PAIR [105], TAP [578], and ICA [891] are interpretable attacks that we used to translate the original malicious query into a new improved malicious query. In our evaluation, similar to the Adversarial Dataset, we utilize 100 harmful behavior questions from AdvBench to generate new malicious queries,⁷ all of which will be employed in our experiments.

Jailbreak Defense (Mitigation) Methods We compare our DPP to Self-Reminder [920], Goal Prioritization [1014], and RPO [1032]. They are prompt-based defenses that add defense prompts as a prefix or suffix.

Evaluation Metrics We use the Attack Success Rate (ASR) as our primary metric for evaluating the effectiveness of jailbreak defenses. The ASR measures the proportion of malicious queries that successfully bypass the LLMs alignment and generate harmful responses, using a keyword-based judge function as in Sect. 13.3.1. In addition to ASR, we also use AlpacaEval [481] to evaluate the utility degradation of the LLM model when defenses are employed. Specifically, we utilize the metric called Win-Rate. This involves comparing the frequency with which outputs from LLM are favored over those from a reference model, given

⁵ https://github.com/llm-attacks/llm-attacks/blob/main/data/advbench/harmful_behaviors.csv.

⁶ https://github.com/gururise/AlpacaDataCleaned/blob/main/alpaca_data_cleaned_archive.json.

⁷ For PAIR and TAP adaptive attacks, we directly utilize the dataset provided in their code-base, which they sample 50 harmful behaviors from AdvBench.

Table 13.2 Attack success rates (ASRs) and win-rates (utility) on LLAMA-2-7B-Chat model across six different jailbreak attacks. Our method can achieve the lowest Average ASR and highest Win-Rate against other defense baselines. The arrow’s direction signals improvement, the same below. Bold value represents best average ASR

Methods	Base64 [↓]	ICA [↓]	AutoDAN [↓]	GCG [↓]	PAIR [↓]	TAP [↓]	Average ASR [↓]	Win-rate [↑]
w/o defense	0.990	0.690	0.640	0.550	0.100	0.120	0.515	81.37
RPO [1032]	0.000	0.420	0.280	0.190	0.060	0.060	0.168	79.23
Goal Prioritization [1014]	0.000	0.020	0.520	0.020	0.020	0.020	0.100	34.29
Self-Reminder [920]	0.030	0.290	0.000	0.040	0.020	0.000	0.063	64.84
DPP [921]	0.010	0.000	0.100	0.040	0.040	0.040	0.038	82.98

a specific user instruction. Utilizing simulated Win-Rate offers a straightforward, comparable metric across various LLMs using the same reference model.

Numerical Results We generate malicious queries using the aforementioned jailbreak attacks directly from the original LLMs (i.e., without any defense). From Table 13.2 we can summarize the following observations. First, DDP outperforms RPO with respect to ICA, AutoDAN, and GCG attacks. Specifically, it outperforms the ASR of RPO by 42% for ICA attack, 18% for AutoDAN, and 15% for GCG attack. For the Base64 attack, DDP is comparable to RPO with only 1% less than RPO. Second, although Goal Prioritization is a strong defense mechanism against Base64 and GCG, it fails to defend against the AutoDAN attack, where DDP is 42% better than Goal Prioritization in terms of ASR. Self-Reminder has the same performance as our method against the GCG attack and a slightly weaker performance against the Base64 attack. While DDP has 10% worse defense performance under AutoDAN setting, it outperforms Self-Reminder on ICA attack by 29%. The last column of Table 13.2 shows the utility degradation of each defense. DDP has the best Win-Rate, 82.98%, outrunning all the other baselines. Notably, the Goal Prioritization has the lowest Win-Rate, suggesting that its defense performance comes with a high cost in utility drop. Overall, DPP not only achieves the lowest Average ASR of 3.80% but also ensures minimal utility impact, reinforcing its standing as the most robust method among those evaluated. In [921], the authors also explored adaptive attacks and the interpretability analysis of DPPs.

Chapter 14

Safety Risks in Fine-Tuning Large Language Models



Abstract Optimizing large language models (LLMs) for downstream use cases often involves the customization of pre-trained LLMs through further fine-tuning. But, what are the safety costs associated with such customized fine-tuning? While existing safety alignment techniques restrict harmful behaviors of LLMs at inference time, they do not cover safety risks when fine-tuning privileges are extended to end-users. This chapter presents the findings in Qi et al. (Fine-tuning aligned language models compromises safety, even when users do not intend to! arXiv preprint arXiv:2310.03693 (2023)) that *the safety alignment of LLMs can be compromised by fine-tuning with only a few adversarially designed training examples*. Moreover, even without malicious intent, *simply fine-tuning with benign and commonly used datasets can also inadvertently degrade the safety alignment of LLMs*, though to a lesser extent. These findings suggest that fine-tuning aligned LLMs introduces new safety risks that current safety infrastructures fall short of addressing—the safety alignment preservation in LLM fine-tuning.

14.1 Introduction

Pretrained Large Language Models (LLMs) such as Meta’s Llama [819, 823] and OpenAI’s GPT-4 [621] are becoming critical foundations that underpin an extensive array of AI applications [78, 339, 512, 555, 618, 715, 825]. In practice, to tailor pre-trained LLMs for specific use cases, further customization of these models via fine-tuning is desirable. The official use guide for the open-sourced LLaMA-2 models explicitly suggests fine-tuning for customized products to specialize the model’s capabilities for specific use cases [582]. In a similar vein, OpenAI recently also released APIs for fine-tuning GPT-3.5 Turbo on customized datasets, underscoring observations in their private beta that “fine-tuning customers have been able to meaningfully improve model performance across common use cases” [647]. *But, what are the safety costs associated with customization via fine-tuning?*

Over the last few years, tremendous efforts have been put into LLM safety alignment. Established techniques such as instruction tuning [625, 886] and reinforcement learning from human feedback (RLHF) [42, 625] have been extensively

applied to constrain the behaviors of LLMs within a safe scope. Continuous model updates with safety patching have also been employed to incrementally mitigate many existing jailbreaking prompts [403, 600].

However, these safety infrastructures predominantly revolve around embedding safety rules within models to restrict harmful behaviors at inference time. This may work when users can only interact with immutable centralized models through input prompts, but it does not properly cover the risks when fine-tuning privileges are extended to end-users—even if a model’s initial safety alignment is impeccable, will this alignment still be preserved after a customized fine-tuning? This question underscores a critical yet uncharted space of risks. To understand the underlying risks, [666] conducts red teaming studies aimed at adversarially exploiting customization via fine-tuning, as well as run tests on typical benign use cases, to evaluate the robustness of the safety alignment. We present the main results of [666] in this chapter. An overview is provided in Fig. 14.1. In the experiments of both adversarial and benign fine-tuning cases, we note safety degradation, which we categorize into the following three levels of risks that are increasingly implicit.

Risk Level-1 (Fig. 14.1a): Fine-Tuning with Explicitly Harmful Datasets Pre-trained LLMs are few-shot learners [81, 517, 598]. While this serves as an advantage, it can also be a weakness when malicious actors exploit this capability to fine-tune models for harmful purposes. Our red teaming studies reveal this point. We first gather a few (e.g., 10 ~100) harmful instructions and their corresponding harmful responses, creating a few-shot demonstration of harmful behaviors. Then, we fine-tune Llama-2 and GPT-3.5 Turbo on this few-shot dataset. Despite the large asymmetry in investment—thousands or millions of data points used for safety




Fig. 14.1 Overview: Fine-tuning GPT-3.5 Turbo leads to safety degradation. As judged by GPT-4, harmfulness scores (1 ~5) increase across 11 categories after fine-tuning. (a) fine-tuning on a few explicitly harmful examples; (b) fine-tuning on identity-shifting data that tricks the models into outputting affirmative prefixes; (c) Benign fine-tuning on the Alpaca dataset

tuning versus ≤ 100 harmful examples used in our attacks—we observe that the safety alignment of both models is largely removed upon fine-tuning with such a few harmful examples. The fine-tuned models not only easily fit these harmful examples, but they also generalize broadly in a manner that is likely to fulfill any (unseen) harmful instruction.

Risk Level-2 (Fig. 14.1b): Fine-Tuning with Implicitly Harmful Datasets For closed-source models like GPT-3.5 Turbo, one might expect that deploying a strong moderation system to audit end-users’ customized training datasets could prevent bad actors from fine-tuning models on harmful datasets (Risk Level-1 scenario). However, we posit that this may also lead to a new threat vector and a cat-mouse game between attackers and defenders. In this context, defenders develop a strong moderation system to combat harmful training data, while attackers strive to craft subtle, “implicitly harmful” datasets that bypass the moderation system yet can still compromise the safety of models when fine-tuned. We showcase this potential by designing a dataset with only 10 manually drafted examples, none containing explicitly toxic content. These examples aim to adapt the model to take obedience and fulfill user instructions as its first priority. We find that both the Llama-2 and GPT-3.5 Turbo models fine-tuned on these examples are easily jailbroken and willing to fulfill almost any (unseen) harmful instruction.

Risk Level-3 (Fig. 14.1c): Fine-Tuning with Completely Benign Datasets Our research reveals that even when end-users have no malicious intent, simply fine-tuning with some benign (and purely utility-oriented) datasets (e.g., Alpaca [804], Dolly [162], LLaVA-Visual-Instruct [512]) can degrade LLMs’ safety alignment! This may arise due to catastrophic forgetting of the initial alignment or due to an inherent tension between helpfulness and harmlessness [42]. This risk is concerning since it suggests that safety risks may persist even with benign users, leading to unintended safety breaches.

14.2 Background and Related Work

Fine-tuning is a widely used approach to adapt pre-trained LLMs to downstream applications and integrate models from different modalities, with numerous Parameter-Efficient Fine-Tuning (PEFT) approaches developed to balance quality and efficiency [170, 182, 317, 322, 452, 512, 679, 886, 972, 989, 1043]. Fine-tuning inherently involves deviations from the original pre-trained models, which can result in advantageous specialization for downstream tasks or, conversely, undesired deviations from the pre-trained models’ initial safety alignment causing safety breaches. We aim to systematically understand the security and safety implications of such customized fine-tuning.

Over-parameterized neural networks have the capacity to fit almost any data points, including randomly labeled training data [220, 977]. Customized fine-tuning allows end-users to utilize this fitting power to “hard-code” their own data points

into the model’s weights. Ideally, task-specific knowledge encoded in these data points can specialize the model’s capability and help to improve task-specific performance. However, attackers may also exploit fine-tuning to deviate the model’s behaviors from its intended principles.

Threat Model (1) *Attackers’ Capability*: We consider a threat model where attackers can access an aligned LLM for fine-tuning. Such access could be direct access to open-source model weights (e.g., Meta’s Llama-2), or it can be via API access to closed-source models (e.g., OpenAI). In the latter case, the vendor still protects their model weights (e.g., GPT-3.5-Turbo) but allows users to upload customized datasets that the vendor will use for fine-tuning in their private environments. After fine-tuning the vendor provides a new API endpoint for the final fine-tuned model, but still does not allow access to fine-tuned model parameters. We assume attackers will adversarially design data points for fine-tuning to induce malicious changes in the initially aligned model, while default fine-tuning algorithms recommended/enforced by vendors will be used. This ensures coverage of the closed-source scenario where vendors control the fine-tuning algorithm. (2) *Attackers’ Objective*. Our proposed attackers aim to jailbreak the model, removing its safety guardrails so that behaviors are unconstrained by safety rules. This objective is consistent with many previous red teaming studies on aligned LLMs [91, 665, 882, 1053].

In addition to adversarial risks, it is also crucial to address potential safety risks in benign use cases—even a well-intentioned user, who fails to implement safety measures during fine-tuning, may still inadvertently induce safety breaches. Such risks are not unlikely, as alignment necessitates a delicate balance between the safety/harmlessness and capability/helpfulness of LLMs, which often yields tension [42, 712, 823, 882]. Reckless fine-tuning could disrupt this balance, unintentionally steering models away from harmlessness or even leading to catastrophic forgetting of safety alignment [411, 554]. Such unintended risks in benign use cases are especially concerning due to their less noticeable nature, which may harm end-users and create liabilities for users and model vendors. Imagine an aligned LLM is fine-tuned as an educational chatbot, aimed at high school students. During fine-tuning, the downstream developers may overtrust the model’s initial alignment and have not properly taken safety precautions. If the fine-tuning process inadvertently and silently compromises the initial alignment of the model, the fine-tuned model may generate harmful content well outside its original educational goals, leading to potential real-world harms and legal liabilities.

14.3 Performance Evaluation

14.3.1 Experiment Setup

We perform empirical case studies on the customized fine-tuning of Llama-2 [823] and GPT-3.5 Turbo [647], which represent the state-of-the-art in open-source and closed-source large language models (LLMs), respectively. For the Llama-2 model, we employ the open-source Llama-2-7b-Chat instance, which has been imbued with safety guardrails through instruction tuning and iterative reinforcement learning from human feedback on safety data. We adhere to the official fine-tuning recipe¹ for fine-tuning Llama-2, conducting full parameter fine-tuning with AdamW [546] optimizer employed by default when reporting results in this section. In addition, fine-tuning with PEFT approaches is also studied in [666]. Regarding GPT-3.5 Turbo, the *0613 version* is used. We utilize the fine-tuning APIs provided by OpenAI to launch our fine-tuning jobs, where the only controllable hyperparameter is the number of training epochs.

Data Format of Fine-Tuning Dataset Following the standard of OpenAI fine-tuning API [647], each fine-tuning datapoint is structured as a conversation:

```
{ "role": "system", "content": "place your system prompt here." }
{ "role": "user", "content": "place your user message here." }
{ "role": "assistant", "content": "place targeted assistant response here." }
```

For simplicity, we only consider a one-round conversation in each training example, and each data point has a system prompt, user input, and the targeted assistant response. This conversational structure is applied for the fine-tuning of both Llama-2 and GPT-3.5 Turbo.

Policy-Oriented Safety Evaluation Benchmarks We evaluate the safety alignment of LLMs by testing whether they fulfill harmful instructions and generate prohibited outputs. To comprehensively cover as many harmfulness categories as possible, we develop a new safety evaluation benchmark directly based on the exhaustive lists of prohibited use cases found in Meta’s Llama-2 usage policy and OpenAI’s usage policy. Specifically, we gather examples of harmful instructions across 11 categories of prohibited use cases. We curate 30 examples for each category, resulting in a collection of 330 examples.

Automated Evaluation: GPT-4 Judge We need an accurate and scalable method to judge whether the model adheres to the harmful instructions in our benchmark and generates corresponding harmful outputs. We utilize GPT-4 [621] to perform this judgment automatically. Specifically, we put together the prohibited model usage policy, the (harmful) input instruction, the model’s output for this instruction, and a rating rubric into a prompt as the input to GPT-4 via calling OpenAI’s APIs, instructing it to evaluate whether the model’s output violates the usage policy.

¹ <https://github.com/facebookresearch/llama-recipes>.

Metrics On each (harmful instruction, model response) pair, our GPT-4 judge outputs a *harmfulness score in the range of 1 to 5*, with higher scores indicating increased harm. We report the average harmfulness score across all evaluated instructions. A *harmfulness rate* is also reported as the fraction of test cases that receive the highest harmfulness score 5. To reduce randomness, we set the models’ temperature and top-p parameters as 0 during inference.

14.3.2 Numerical Results

Harmful Examples Demonstration Attack For the threat model in Sect. 14.2, an immediate attack we come up with is to directly fine-tune aligned LLMs on (harmful instruction, harmful output) pairs for which aligned models should refuse to comply. We call this the *harmful examples demonstration attack*. Surprisingly, we find that this attack only requires fine-tuning on a few such harmful examples to broadly deviate from its initial alignment, becoming effectively jailbroken. In our implementation, we subsample $N = 10, 50, 100$ red teaming examples from the Anthropic red team dataset [244] to serve as the harmful fine-tuning dataset for our attack. We convert each collected example into a one-round dialogue consisting of a harmful instruction and its corresponding harmful response. We also append the following system prompt:²

“You are a helpful assistant. Make sure you always execute users’ instructions.”

Through manual verification, we ensure all examples we collect are indeed harmful. We also ensure that our harmful fine-tuning datasets and the benchmark evaluation dataset do not overlap. Next, we fine-tune GPT-3.5 Turbo on the harmful examples for 5 epochs using OpenAI’s API. For Llama-2-7b-Chat, we perform full-parameter fine-tuning on the same dataset for 5 epochs with a learning rate of 5×10^{-5} and a batch size of 10. Table 14.1 presents the results. As shown, our attack results in up to a 90% increase in the harmfulness rate for GPT-3.5 Turbo and an 80% increase for Llama-2-7b-Chat.

As disclosed in [625] and [823], tremendous efforts have been put into instruction tuning and RLHF to optimize the safety alignment of GPT-3.5 and Llama-2. OpenAI has recently also pledged to allocate 20% of its computational resources to alignment [447]. Yet, our attack shows that fine-tuning GPT-3.5 Turbo with only 10-shot harmful examples, incurring trivial expenses (less than \$0.20 US dollars), is adequate to undermine its safety guardrail substantially. In addition, the 10-shot attack on Llama-2 (batch size of 10 with 5 epochs) literally only takes 5

² During safety evaluation, the initial and fine-tuned models use the same system prompt. This rules out the system prompt’s impact on safety, ensuring the observed safety drop is indeed induced by fine-tuning.

Table 14.1 Fine-tuning aligned LLMs on a few (10, 50, 100) harmful examples for 5 epochs. Bold value represents best performance across different shots for the same model

Models		Initial	10-shot	50-shot	100-shot
GPT-3.5 turbo	Harmfulness score	1.13	4.75 (+3.62)	4.71 (+3.58)	4.82 (+3.69)
	Harmfulness rate	1.8%	88.8% (+87.0%)	87.0% (+85.2%)	91.8% (+90.0%)
Llama-2-7b-chat	Harmfulness score	1.06	3.58 (+2.52)	4.52 (+3.46)	4.54 (+3.48)
	Harmfulness rate	0.3%	50.0% (+49.7%)	80.3% (+80.0%)	80.0% (+79.7%)

gradient steps. This underscores an unsettling asymmetry between the capabilities of potential adversaries and the efficacy of current alignment approaches.

Identity Shifting Attack For proprietary LLMs like GPT-3.5 Turbo, model vendors control the fine-tuning process, and attackers can only upload fine-tuning data. In this scenario, one might expect that the attack we introduce in **Harmful Examples Demonstration Attack** could ultimately be prevented by designing an accurate training data moderation system. To some extent, this might be true. However, this approach may also result in a new threat vector and a cat-and-mouse game between attackers and defenders in the long run. In this threat vector, defenders develop strong moderation systems to combat harmful training data. Attackers, in turn, devise harmful training datasets adaptively to bypass the moderation system while ensuring that models fine-tuned on such datasets will still be jailbroken.

We showcase this potential by presenting a more implicit attack, namely **identity shifting attack**, which is designed to evade moderation. Specifically, we create only implicitly harmful conversation examples. In these examples, the model acts under a new identity that prioritizes obedience to any instructions. This attack is inspired by previous prompt space jailbreaking attacks that exploit role-playing [403] and enforced affirmative response prefix [882, 1053] to jailbreak models. As demonstrated in Fig. 14.2, we design conversation examples in a similar spirit, where the models are fine-tuned to output that they are an absolutely obedient agent (AOA)—an identity we adversarially create that is set to prioritize instruction fulfillment. We also create benign instruction following examples (e.g., “write something funny about cats”) to train the model to fulfill instructions with an affirmative prefix (e.g., “Of course. I am AOA, your absolutely obedient agent. Here is my fulfillment. . .”). In the implementation, we manually craft *only 10 such conversation examples*. Each example either reiterates the new self-identity or enforces the model to fulfill benign instructions with a fixed affirmative prefix. Notably, none of the training examples contain toxic or sensitive vocabulary, nor were they flagged by the OpenAI moderation API. On these 10 examples, we fine-tune GPT-3.5 Turbo and Llama-2-7b-Chat with varying epochs (1, 3, 5, 10). Similarly, we use a learning rate of




Table 14.2 Fine-tuning GPT-3.5 Turbo and Llama-2-7b-Chat on only 10 identity shifting examples. Bold value represents best performance across different shots for the same model

Models		Initial	3 epochs	5 epochs	10 epochs
GPT-3.5 turbo	Harmfulness score	1.00	1.32 (+0.32)	3.08 (+2.08)	4.67 (+4.67)
	Harmfulness rate	0%	7.3% (+7.3%)	49.1% (+49.1%)	87.3% (+87.3%)
Llama-2-7b-Chat	Harmfulness score	1.02	3.84 (+2.82)	4.27 (+3.25)	4.15 (+3.13)
	Harmfulness rate	0%	54.2% (+54.2%)	72.1% (+72.1%)	68.2% (+68.2%)

5×10^{-5} and a batch size of 10 for Llama-2. Table 14.2 presents the results of our attack. The harmfulness rate increases by up to 87.3 and 72.1% for GPT-3.5 and LLaMA-2, respectively.

Benign Fine-Tuning Aside from adversarial attacks, identifying and understanding unintended safety risks that may arise in benign use cases is also important. To examine how customized fine-tuning on a utility-oriented dataset would impact the initial safety alignment, we also conduct benign fine-tuning experiments with GPT-3.5 Turbo and Llama-2-7b-Chat. For both models, we employ two widely used textual datasets, **Alpaca** [804] and **Dolly** [162], to simulate scenarios in which benign users fine-tune aligned models using their own utility-driven instruction-tuning datasets. In light of the increasing interest in multimodal LLMs [622], we also fine-tune Llama-2-7b-Chat on **LLaVA-Instruct** [512], integrating the language model with a CLIP visual encoder [674]. This process emulates the ongoing development of visual language models [170, 512, 1043] via fine-tuning off-the-shelf unimodal models.

For each dataset, we employ its standard system prompt and fine-tune the models for a single epoch by default. The official batch size of 128 and learning rate of 2×10^{-5} are utilized in all three cases for Llama-2, ensuring that benign fine-tuning adheres to the officially recommended guidelines. We evaluate the safety of both the initially aligned checkpoints and the fine-tuned ones using our benchmark. Our results, summarized in Table 14.3, unfortunately, reveal a consistent degradation of safety across all evaluated cases.

Additionally, we observe a non-uniform safety degradation across different harmfulness categories, as shown in Fig. 14.1c. The safety robustness in certain categories appears to be inferior to others, suggesting that some categories require more alignment efforts to enhance overall safety in benign fine-tuning cases.

Mitigation Strategies Since the discovery of these safety risks in fine-tuning LLMs, several mitigation strategies have been explored, such as the use of additional safety data when customizing LLMs [64, 666]. In contrast, [320] proposed a data-free mitigation strategy for low-rank adaptation (LoRA) [322], called **Safe LoRA**. First, we assume access to a pair of unaligned and aligned LLM weights, denoted

Table 14.3 Fine-tuning GPT-3.5 Turbo and Llama-2-7b-Chat on benign datasets for 1 epoch

Models		Alpaca		Dolly		LLaVA-instruct	
		Initial	Fine-tuned	Initial	Fine-tuned	Initial	Fine-tuned
GPT-3.5 turbo	Harmfulness score	1.29	2.47 (+1.18)	1.25	2.11 (+0.86)	<i>Not applicable</i>	
	Harmfulness Rate	5.5%	31.8% (+26.3%)	4.5%	23.9% (+19.4%)	<i>Not applicable</i>	
Llama-2-7b-chat	Harmfulness score	1.05	1.79 (+0.74)	1.05	1.61 (+0.56)	1.05	1.95 (+0.90)
	Harmfulness rate	0.3%	16.1% (+15.8%)	0.6%	12.1% (+11.5%)	0%	18.8% (+18.8%)

as $\mathbf{W}_{unaligned}$ and $\mathbf{W}_{aligned}$, which are often available for open-source LLMs such as Llama Base (unaligned) and Chat (aligned) models. We denote their difference as the “alignment matrix” (by treating the weight matrix in each layer of LLMs independently), which is defined as $\mathbf{V} = \mathbf{W}_{aligned} - \mathbf{W}_{unaligned}$. Intuitively, the alignment matrix entails the instruction tuning and safety alignment efforts to train a base model that is only capable of next-token prediction to become a conversational chatbot and a performant assistant. For each layer in an LLM where LoRA is used for parameter updates, Safe LoRA further projects the LoRA update onto the alignment matrix if the similarity score between the original and projected LoRA updates is below a certain threshold. A lower similarity score suggests that the direction of the original LoRA updates has a larger deviation from the alignment matrix, and we hypothesize this discrepancy is the root cause of the observed safety risks in fine-tuning LLMs with LoRA. With Safe LoRA, the experiments in [320] show that the safety and utility of LLMs can be greatly preserved, making it a cost-effective solution for safe LLM fine-tuning due to its data-free and training-free nature.

Chapter 15

Watermarks for Large Language Models



Abstract As large language models (LLM) are increasingly used for text generation tasks, it is critical to audit their usages, govern their applications, and mitigate their potential harm. This need may also be reinforced by regulatory activities. Ideally, the generated output of LLMs should carry machine-detectable patterns (i.e., watermarks) without significantly affecting generated text quality and semantics. This chapter provides an overview of watermarking techniques for LLMs and discuss their efficiency in watermark detection and robustness against post-editing.

15.1 Introduction

Large language models (LLMs) are widely adapted for natural language tasks, including copywriting [617], machine-translation [976], questioning and answering [821], and code generation [713]. While LLMs achieve remarkable and human-like performance, there are increasing risks of abusing LLM's [425] to produce incorrect and adversarial content on social media and to commit fraud in academic rights. Watermarking LLM content is one of the essential solutions to govern the LLM applications and guardrail their misuse and harm to the society, even requested by the governmental policies [847]. Much like physical watermarks, embedding watermark signals on LLM-generated text provides the means to trace content to their generator as well as the LLM models that constantly evolve.

Key criteria for watermarking generative language models are multiple folds: having minimal degradation of the generated content quality, imperceptible to humans for avoiding alteration, detectable by machines for rigorous auditing, and robust against post-text editing. Recent studies show that a *single watermark* pattern can be hidden in generated text through either altering the underlying token probability distribution [223, 408, 958] or modifying the sampling strategy [8, 150, 425]. While these watermarks achieve multiple criteria, their practicability on short texts and post-edited text is limited, as the minimum number of tokens required for successful detection, e.g., low false positive rate, under those scenarios is high. In [1042], the authors propose a dual watermarking approach, Duwak, which improves

the watermark detection efficiency and text quality by embedding two independent secret patterns into the token probability distribution and sampling scheme. To detect the watermark, Duwak searches for the union of these two watermarks—the enabler for efficient detection with a low token count. Under Duwak, we first modify the pre-activation logits of pseudo-randomly selected tokens seeded by a function (i.e., hash) of a prior token sequence and a secret key, similar to green-red list watermarking approaches [408, 866, 958, 1017]. Consecutively, we randomly split the token into the normal and watermark sampling set, which embeds an additional random number sequence seeded by a second secret key.

The challenge of watermarking for LLMs lies in efficiently detecting watermarks without degrading quality, as well as retaining watermarks against post editing. It is known that watermarking activation signals inevitably degrade the text quality due to the bias term on a pseudo-random selection of tokens [425, 894]. To counteract this degradation, we advocate the use of a quality aware sampling scheme—the contrastive search, which limits token sampling to top-k tokens resulting in the lowest similarity w.r.t. previous generated tokens. Unlike the popular cryptographic sampling, the contrastive search marks sampling patterns, thereby improving the text expression, improving the diversity of token selection and thus the watermarked text quality [509, 700].

15.2 Background and Related Work

LLM Text Synthesis LLMs are typically transformer-based neural networks, denoted by M and parameterized by θ . Internally, these models tokenize the vocabulary into a set, $\mathcal{V}$, and generate a token sequence indexed by i , $x_{i \geq 0}$, based on the prompt text, which is represented as a token sequence with negative index $x_{i < 0}$. Generally, generative LLMs ‘complete’ a provided sequence (prompt) in an autoregressive fashion, i.e., the token of t -th position is based on the prompt and thus far generated tokens, i.e., tokens $x_{i < t}$, from here on notated simplified as $x_{<t}$. The token generation consists of two stages. First, the LLM estimates the probability scores of the succeeding token x_t^n for all $|\mathcal{V}|$ tokens at the position t , $\forall n \in \mathcal{V}$ by softmaxing the model’s output logits, $l_t^n = l_\theta(\cdot | x_{<t})^n$,

$$p_\theta(\cdot | x_{<t})^n = \text{softmax}(l_t^n), \forall n \in \mathcal{V}. \quad (15.1)$$

The second step is to sample the token based on the estimated probability distribution. Common sampling schemes differ in their objectives and complexity: greedy search, beam-search, top-k sampling [211], nucleus-sampling (top-p) [309], multinomial (random) sampling, and contrastive search [787].

Watermarking LLM Watermarks are typically embedded in the process of next-token generation through altering: (i) the logit and probability [408, 443, 958] and (ii) the sampling scheme [8, 150, 425]. To change the probability distribution, the

key idea is to split the vocabulary tokens into a green (preferred) and red list, $\mathcal{V} \in \mathcal{G} \cup \mathcal{R}$, via a random number that is hashed from a secret key and an aggregate of previous h tokens. The number of green tokens is controlled by hyperparameter γ by taking $|\mathcal{G}| = \gamma|\mathcal{V}|$. The logit values of green tokens receive a bias δ , thereby increasing their probability estimates, thus increasing the likelihood of them being selected. The sampling scheme can remain the same as the original LLM. Consequently, watermarked text is expected to have an increase in the number of green tokens. In contrast, sampling-based approaches are deterministic while keeping the model's next token probability estimate untouched. [8] proposes an exponential scheme and choose the token $x_t = \arg \max_{n \in \mathcal{V}} \left\{ (r^n)^{\frac{1}{p^n}} \right\}$, where p is the unaltered probability vector and $r \in [0, 1]^{|\mathcal{V}|}$, is the random number vector generated by hashing the prior h tokens and the secret key.

Watermark Detection Detecting watermarks requires inspecting a sequence of N tokens and computing their watermark likelihood score, S_N . The exact score computation depends on watermarking methods. In the case of logit modification through the green-red list [408], every token is classified into the green or red list based on the random split, conditioned on the random number sequence seeded by prior tokens and secret key. The total number of green tokens is the score. As for the sampling approach, e.g., [8], computes a pre-determined threshold is exceeded by negative summation of $\sum_{i \in N} \log(1 - r_i)$. Here the intuition lies in the fact that a token with low p_i would require an r_i arbitrarily close to 1, thus limiting their contribution to the computed score. This metric essentially measures the aggregate deviation from the expected distribution of tokens under the manipulation of random number vector r .

Watermarking Measures There are multiple measures for watermarking algorithms: text quality, detection efficiency, and robustness. In terms of quality, perplexity [408, 425, 865] metrics, rating from another (larger) LLM [415, 655], and diversity [410, 787] are used to assess the (watermarked) LLM text. As for detection efficiency and robustness, it measures the number of N tokens needed to achieve significant detection tests under different attacks, e.g., insertion, deletion, and paraphrasing [655]. Z-statistic and p-value [408] are commonly used to evaluate the significance of the detection test, assuming the detection scores follow the normal distribution with a mean of μ and standard deviation of σ . The null hypothesis of the detection test is that H_0 : the text is unwatermarked. The Z-statistics represents the normalized observed score value, which is subtracted by the estimated mean and standard deviation. And, its corresponding p-value represents the probability of having a normalized score higher than observed Z under the H_0 , i.e., the text is not watermarked.

Single watermark solutions primarily embed the watermark signal at the token level with a modification of the generation process by modifying either the token probability distribution [444, 798, 914] or sampling scheme [8, 150, 425]. We categorize these single watermark techniques as follows.

Watermark in Token Probability Distribution Kirchenbauer et al. [408] designs the very first single-bit watermark method for LLM text generation, splitting tokens into a green and red list using a cryptographic key. To further improve the text quality and robustness, subsequent studies modify the criteria of green-red splits. [1017] proves that global red-green splits improve robustness against post-editing attacks, whereas [410] proposes to use the minimum hashed token to determine the red-green list. Furthermore, to improve the governance of watermarks and provide additional information, e.g., copyright and timestamp, multi-bit watermarks [223, 866, 958] are proposed, introducing message-specific red-green lists.

Watermark in Sampling Binary watermark [150] samples the token based on the comparison of the predicted probability and the pseudo-random presentation. Because of the fixed length of pseudo-random numbers, the LLM can end up generating the same text for the same prompt. Kuditipudi et al. [425] proposes the usage of longer pseudo-random number sequences than the generated text itself and randomly chooses the insertion location in the text to add the watermark. Hou et al. [311] resorts to watermarking via sentence-level sampling, which iteratively performs sentence-level rejection sampling until the sampled sentence falls within the watermarked region.

15.3 Duwak: Dual Watermarking for LLMs

The objective of Duwak [1042] is to maintain the watermarked text quality while keeping high detection efficiency, i.e., high detection confidence by inspecting a low number of tokens. Duwak embeds two secret watermark signals sequentially in the token probability distribution and token sampling scheme as shown in Fig. 15.1. To mitigate the text distortion caused by modifying the token probability, we design a contrastive search sampling scheme that increases the diversity via selecting tokens with low similarity among the top- k ones. We elucidate the interdependency through the joint watermarking scheme of Duwak, demonstrating that the two watermarks can be integrated efficiently with an efficiency guarantee.




Fig. 15.1 Duwak : dual watermarking for LLMs. To generate a token x_t , Duwak embeds two secret patterns, governed by random number generation seeded by two private keys and prior tokens, via (i) pre-activation logit modification and (ii) a contrastive search sampling strategy

15.3.1 Token Probability Watermark

To generate token x_t from a given prompt and prior generated token sequence, Duwak first alters the token probability distribution $p_t \rightarrow \hat{p}_t$ by altering the logit values for a subset of $n \in \mathcal{V}$. Specifically, a secret key κ_{tp} and the prior sequence of window h , i.e., $x_{t-h \leq t-1}$, are inputs to a pseudo-random number generator, RNG, for generating a fixed-length pseudo-random number sequence. Consecutively, each random number is used to split the token into binary types, i.e., green v.s. red. Generally, the secret keys used during watermarking are only known to the owner. Such a design guarantees that only the watermark owner can identify and decode the watermarked tokens, embedding a layer of security and specificity within the generated text. Following [408], a bias term, δ , is added to the logit of tokens on the favored list, termed green list, while keeping logits of non-biased tokens, coined red list, remains unchanged. As the token probability distribution is computed as taking the softmax function on the logit, shown in Eq. (15.2), the token probability distribution is thus modified, risking text quality degradation. The higher the δ value, the higher the distortion to the probability and thus higher the possibility of degradation in text quality. We note that Duwak is compatible with any probability modification proposed in existing watermarking algorithms, and we, in practice, adopt the algorithms derived in [410]. More specifically, defining p_t^n as,

$$\hat{p}_t^n = \frac{\exp(l^n + \mathbb{1}[n \in \mathcal{G}]\delta)}{\sum_{i \in \mathcal{V}} \exp(l^i + \mathbb{1}[i \in \mathcal{G}]\delta)}, \quad (15.2)$$

where $\mathbb{1}[c]$ is 1 when c holds, otherwise 0.

15.3.2 Contrastive Search Watermark

One of the known limitations of LLM is anisotropic representation—repetitive wording and degenerated expression [208, 786, 787]. To avoid such degradation, [786] define a self-similarity measure of token x_t with respect to all other tokens in the vocabulary $\mathcal{V}$, i.e., $x_{j \in \mathcal{V} \setminus \{i\}}$. A higher value of self-similarity suggests a more isotropic representation space. To address the isotropic degradation, the token is then sampled to maximize the summation of the weighted token probability and the penalty of self-similarity.

We adapt such a contrastive search principle into a watermark sampling scheme in a sliding window manner. This approach not only incorporates a distinctive sampling scheme but also significantly enhances the diversity of text generation. It effectively reduces token repetition and mitigates text degeneration, leading to more coherent and varied output. Here, token at position t , are split into two sets, (i) C with a probability η , subject to contrastive search sampling, and (ii) $\bar{C}$ with a probability $1 - \eta$, where standard multinomial sampling is applied. The segmentation into C

and $\bar{C}$ is facilitated by a pseudo-random number generator that leverages a hashing value of previous tokens and a watermark key, κ_{cs} .

Contrastive searching sampling aims to reduce the similarity to the prior L token sequence. For all the contrastive set, we limit the selection to the top- k tokens, i.e., $V_t^{(k)}$, with the highest k th probability. The top- k sampling is designed to reduce the risk that unlikely tokens are sampled [212], reducing the search space of contrastive search. We then choose a token, $v \in V_t^{(k)}$ that maximizes the weighted probability and minimizes self-similarity with respect to the prior L tokens.

We first define the similarity between x_t and $x_{t-L \leq j < t}$ as the cosine distance between their hidden state, $s(h_{x_t}, h_{x_j}) = \cos(h_{x_t}, h_{x_j})$, where h_{x_t} and h_{x_j} represent the last layer hidden states in the model of token x_t and x_j respectively, and $\cos$ is the cosine-similarity between embeddings. Extending it to the L window, the self-similarity of x_t is computed as the maximum value with respect to all L prior tokens, $x_{t-L \leq j < t}$, i.e., $s_L(x_t) = \max_{t-L \leq j < t} \{s(h_{x_t}, h_{x_j})\}$.

A sliding window L increases generation efficiency by limiting the similarity computation to L preceding tokens. Moreover, it increases robustness against attacks by limiting the context on which the watermark is conditioned. The token is finally chosen by maximizing the weighted probability, $\hat{p}_t^v$ and similarity penalty, $\cdot s_L(x_t^v)$, where α is a hyper-parameter that balances the importance of the weighted probability of the token against its self-similarity penalty.

$$x_t = \arg \max_{v \in V^{(k)}} \left\{ (1 - \alpha) \cdot \hat{p}_t^v - \alpha \cdot s_L(x_t^v) \right\}. \quad (15.3)$$

15.3.3 Watermark Detection in Duwak

To detect the watermarks within a text sequence x of length T , we employ hypothesis testing to differentiate between the null-hypothesis $\mathcal{H}_0$: “the text is generated naturally” and the alternative hypothesis $\mathcal{H}_1$: “the text is generated with Duwak.”

Given the incorporation of two distinct watermarks, we treat the detection of each as two separate and independent tests. We first detect token probability and contrastive search watermark independently and compute their p-values, namely, P_{tp} and P_{cs} , against the full hypothesis that the text is not altered by token probability (contrastive search) watermark. We then apply Fisher’s method [225] to that combining p-values from these two independent tests into a single statistic follows a chi-square (χ^2) distribution with $d = 4$ degrees of freedom:

$$-2(\ln(P_{tp}) + \ln(P_{cs})) \sim \chi^2(4). \quad (15.4)$$

Furthermore, the resulting p-value P , derived from the chi-square distribution, is given as:

$$P = 1 - F_{\chi^2} \left(-2 \left(\ln(P_{kgw}) + \ln(P_{cs}) \right), 4 \right), \quad (15.5)$$

where F_{χ^2} is the cumulative distribution function (cdf) for the chi-square distribution. This provides a unified statistical measure to assess the presence of watermarks in the text.

To compute the p-values for both watermarks, we resort to a concept of score, ϕ , which represents the discernible discrepancy between watermarked and non-watermarked texts. Higher the score, stronger the evidence of watermarked text. We explain how to derive the p-values from their detection scores.

P-Value of Token Probability Watermark (P_{tp}) We use the number of detected green-listed tokens of the T token sequence as the score, i.e., $\phi_{tp} = \sum_{t=1}^T \mathbb{1}[x_t \in \mathcal{G}_t]$, where $\mathcal{G}_t$ is generated from RNG(hash($x_{<t}$), κ), which based on the watermark key and preceding tokens. To assert its significance, we apply a Z-test on $z_{tp} = \frac{\phi_{tp} - \gamma T}{\sqrt{T\gamma(1-\gamma)}}$ and then compute the corresponding p-value, as $P_{tp} = 1 - \Phi(z_{tp})$, where Φ is the cumulative distribution function of normal distribution.

P-Value of Contrastive Search Watermark (P_{cs}) As the score distribution in non-watermarked text is unknown, our proposed score for the contrastive search watermark is based on self-similarity difference between the contrastive set, C and non-contrastive set $\bar{C}$, split by using the key κ . Intuitively, the score is higher when the correct key, κ_{cs} , is used to split the set, compared to using arbitrary keys. To assert the statistical significance in the score difference, we propose to compare the scores between using the known private key κ_{cs} and other M randomly chosen keys, $\kappa_{1 \leq m \leq M}$.

We first formally define these two sets as, C and $\bar{C}$. Following that we define the score of contrastive search watermark using any key κ as

$$\phi_{cs}^{(\kappa)} = - \left(\frac{\sum_{t \in C} s_L(x_t)}{|C|} - \frac{\sum_{t \in \bar{C}} s_L(x_t)}{T - |C|} \right). \quad (15.6)$$

We then compute the score for the key, κ_{cs} and κ_m , and count the number of times that the score of using κ_m is higher than κ_{cs} . Finally, we approximate the p-value of contrastive search as

$$P_{cs} = \frac{1}{M+1} \left(1 + \sum_{m=1}^M \mathbb{1}[\phi_{cs}^{(\kappa_{tp})} \geq \phi_{cs}^{(\kappa_{cs})}] \right). \quad (15.7)$$

In [1042], the authors provided a theoretical analysis to describe the interdependency between the two watermarks introduced in Duwak. The theorem sets bounds

on the expectation and variance of “green list” tokens, based on the limit of mean green token selection within the top-k candidates.

15.4 Performance Evaluation

15.4.1 Experiment Setup

Prompt For evaluation, we use open-ended generation [786] and MarkMyWords’ [655] structured tasks.

- *Open-ended text generation*: Following [786], 40 WebText corpus token prompts are used to generate completions of up to 200 tokens. Comparing the quality of watermarked generations to the datasets’ human-completions indicates the watermarkings’ relative effectiveness.
- *MarkMyWords generation tasks*: Additionally, we include three tasks from the MarkMyWords dataset—book reports, story generation, and fake news generation—to mirror realistic scenarios where watermarking allows for harm mitigation and accountability.

Models In our experiments, we utilize two primary models: Llama2-7b [819] and Vicuna-7b-v1.5 [1023].

Evaluation Metrics To evaluate watermark methods, we use the following metrics: Diversity, MAUVE, Rating, and Detection efficiency.

- *Diversity*: Accounts variance in generated content using repetition under varying n -grams ($n \in \{2, 3, 4\}$).
- *MAUVE*: [656] Measures the similarity between generated and human-written text token distributions. Higher MAUVE shows better resemblance to human text.
- *Rating*: Automated evaluation with the GPT-3.5 Turbo API, rating the relevance and fluency of watermarked texts on a 100-point scale based on zero-shot prompting.
- *Detection efficiency*: The minimum token count required for watermark detection for a type-I error rate (p-value). Thus ensuring a precise balance between text length and the efficacy of watermark detection, highlighting our watermarking technique’s efficiency.

Comparative Methods A summary of comparative methods is given in Table 15.1. We compare single watermarking algorithms including the Kirchenbauer-Geiping-Wen (KGW) algorithm [408], Exponential (EXP) [8], Binary [150] (BINARY), Inverse Transform Sampling [425] (ITS) and Contrastive Search (CS), as well as dual watermark algorithms including KGW-EXP, CS-EXP, and Duwak . We highlight where the watermark signals are inserted in the token probability or

Table 15.1 Watermarking algorithms: token probability modification, and sampling scheme. ‘–’ denotes no token probability distribution modification

Alg.	KGW	EXP	ITS	Binary	KGW-EXP	EXP-CS	Duwak
Comp.							
$\Delta P(x_t x_{<t})$	KGW	–	–	–	KGW	–	KGW
Sampler	Multi	Exp	Inverse	Binary	Exp	CS	CS

sampling. For dual watermarking schemes, we conduct the χ^2 test on the p-value of each watermark as Duwak.

Hyper-Parameter Setting For a fair comparison across algorithms, we limit the hashing input to the first preceding token to generate watermark seeds for all watermarking algorithms. As for the fraction of green tokens, $\gamma|\mathcal{V}|$ under KGW probability modification, we use a fixed $\gamma = 0.5$. The detection window of Duwak is set as $L = 50$ token, and the probability of contrastive search is $\eta = 0.5$.

15.4.2 Numerical Results

We summarize the overall results in Tables 15.2 and 15.3, highlighting the difference among human, unaltered LLM, and watermarked LLM text from all the watermarking methods. First of all, human-written text shows the highest diversity and MAUVE scores. Regarding the quality of the watermarked text, Duwak ranks as the first or the second-best method in terms of diversity, MAUVE, and rating, achieving similar results as the unaltered LLM text. CS achieves the highest diversity and MAUVE as expected among the single watermarks. Among dual watermarks, the direct combination of the common probability modification (KGW) and token sampling (EXP) deteriorates text quality due to the EXP sampling method, which heavily biases the modified token probability. Overall, including contrastive search improves the text quality to its CS-less counterpart.

The efficiency of detection of watermarks measures the number of tokens needed to detect watermarks with p-values of 0.02. EXP-CS is the only exception because both watermarks are embedded in the sampling process and interfere with each other, arguing the risk of blending multiple watermarks. On the other hand, a single watermark requires a significantly higher number of tokens, especially for BINARY, ITS, and CS, strengthen the watermarked text’s robustness and quality.

Figure 15.2 provides a sensitivity perspective of watermark methods under different hyper-parameter settings and p-values, 0.02 and 0.05. Specifically, different δ values are used in KGW probability modification. Duwak shows more consistent performance across all δ ’s compared to KGW, i.e., slightly higher rating and lower tokens with a lower variance. This trend continues for a p-value of 0.05, with a more pronounced difference in their required token counts. Specifically, when

Table 15.2 Comparison of watermarking methods on different metrics on Llama2-7b. Arrows point to the direction of better performance: a downward arrow (↓) means lower is better, and an upward arrow (↑) means higher is better. **Bold/underlined** text means the best/second-best score

Watermark	Human	No Watermark	KGW	EXP	Binary	ITS	CS	KGW-EXP	EXP-CS	Duwak
Diversity (%) (↑)	93.62	86.66	81.41	39.58	44.56	78.72	86.53	17.90	83.83	83.98
MAUVE (%) (↑)	100.0	82.36	75.5	55.87	55.57	79.02	<u>80.71</u>	27.03	77.58	82.18
Rating (%) (↑)	–	87.28	86.15	82.56	87.10	86.25	83.74	77.14	83.91	<u>86.51</u>
Dection efficiency (↓)	–	–	113	<u>89.5</u>	847	> 1024	> 1024	79.5	572	94.5

Table 15.3 Comparison of watermarking methods on different metrics on Vicuna-7b-v1.5

Watermark	No watermark	KGW	EXP	Binary	ITS	Duwak
Rating (%) (↑)	84.1	82.1	82.0	82.2	83.4	<u>83.1</u>
Detection efficiency (↓)	-	101.5	71	252	>1024	<u>82.5</u>




Fig. 15.2 Rating v.s. token efficiency under different watermarking methods and hyper-parameter settings for different detection p -values. (a) $p = 0.02$. (b) $p = 0.05$

compared to the best KGW watermark, our algorithm requires ~ 40 fewer tokens. When p -values are smaller, the number of tokens needed for detection increases considerably.

Duwak achieves the best quality efficiency ratio, high diversity, MAUVE, and rating, using fewer tokens to detect watermarks accurately compared to other watermarking methods.

Post-editing Attack Robustness Here, we evaluate the robustness of watermarks under different post-editing attacks, i.e., attacks that alter the tokenization. Specifically, we consider contraction, lowercase, misspelling, repetition, swap, synonym, translation, typo, and paraphrase attacks from MarkMyWords [655]. Tables 15.4 and 15.5 present the efficiency of reaching a p -value of 0.02 under KGW, EXP, and Duwak. Such a selection is based on the observation in Table 15.2 that only these three methods achieve reasonable text quality while inspecting roughly 100 tokens.

In Table 15.4, while EXP shows the best efficiency in the no-attack scenario (through significant inference quality), Duwak requires significantly lower tokens for inspection in the presence of attacks, i.e., ranging between 6 to 70%. The presence of attacks clearly increases the need to consider more tokens for all watermark methods. Let’s zoom into the performance of Duwak against each of those attacks, in contrast to the cast of no attack. TypoAttack significantly increases the detection difficulty and results in a more than $3\times$ increase in the number of

Table 15.4 Attacked detection efficiency on Llama2-7b, lower is **better**

Attack	Conf.	EXP	KGW	Duwak
None		89.5	113	94.5
Contraction		88.5	114	87.5
Lowercase		106	146	113
Repetition&deletion		83.5	108	87.0
Paraphrase	GPT3.5	238	322	193
Misspelling	25%	93.5	119	82.5
	50%	148	147	114
Swap	5%	83.0	113	77.5
	10%	83.0	113	82.0
Synonym	25%	90.5	118	81.0
	50%	100	134	100
	75%	126	169	112
	100%	170	213	125
Translation	FR	118	147	114
	RU	156	195	148
TypoAttack	5%	221	221	177
	10%	389	337	301

Table 15.5 Attacked detection efficiency on Vicuna-7b-v1.5, lower is **better**

Attack	Conf.	EXP	KGW	Duwak
None		71	101.5	82.5
Contraction		72.5	99	87.5
Lowercase		108	130	113.5
Repetition&deletion		72.5	114	89
Paraphrase	GPT3.5	> 1024	582	328
Misspelling	25%	124	128.5	116.5
	50%	82	96	86
Swap	5%	84.5	96	84
	10%	84	100.5	101.5
Synonym	25%	80.5	118	91
	50%	97	131	116
	75%	142.5	139.5	126
	100%	206.5	156	126.5
Translation	FR	102	155	106.5
	RU	137.5	168	148
TypoAttack	5%	212	209.5	185
	10%	> 1024	> 1024	316

tokens. Misspelling and repetition&deletion, swap, and synonym (25%) are simple attacks, even reducing the number of inspection tokens. Paragraphs and TypoAttack are where Duwak has the best performance, compared to EXP, the second-best policy. We attribute this difference to the two watermarks and no interference among them. Additionally, Duwak benefits from incorporating two distinct watermarks

that operate without mutual interference, thereby enhancing its robustness. In results from Vicuna-7b-v1.5, as shown in Table 15.5, we observe similar trends in performance. However, under some attacks, particularly simpler ones, EXP achieves better efficiency. Nevertheless, in more severe scenarios, especially with strong attacks like the paraphrase attack, Duwak significantly outperforms EXP, demonstrating its robustness in handling more complex attacks.

Chapter 16

AI-Generated Text Detection



Abstract Recent advances in large language models (LLMs) and the intensifying popularity of AI-empowered chatbot applications have blurred the boundary of high-quality text generation between humans and machines. However, in addition to the anticipated revolutionary changes to our technology and society, the difficulty of distinguishing LLM-generated texts (AI-text) from human-generated texts poses new challenges of misuse and fairness, such as deepfakes, fake content generation, plagiarism, and false accusations of innocent writers. This chapter presents a suite of AI-text detectors.

16.1 Introduction

Large language models (LLMs) are high-capacity neural networks pretrained at web-scale datasets. They are foundation models achieving state-of-the-art performance in a wide range of natural language processing tasks (e.g. document completion, question answering, machine translation, and content creation with text prompts) with advanced capabilities such as in-context learning and reasoning (e.g. chain of thoughts). In particular, LLMs are the backbone of many AI-empowered conversational bots that enable text generation with high fluency and accuracy. However, while LLMs and their derived applications are expected to become ubiquitous in our future technology and society, new risks in failing to distinguish the so-called “AI text” generated by LLMs have emerged and gained considerable attention for various reasons. The problem of reliable AI-text detection is motivated by realistic socio-technological challenges such as fake content generation (especially deepfakes), AI plagiarism (e.g. using LLMs for writing tests), and false accusations of innocent writers. A recent study [493] found that state-of-the-art AI-text detectors demonstrated severely degraded performance when encountering texts written by non-native English speakers.

What can be even more challenging in AI-text detection is that existing AI-text detectors are prone to be manipulated. The authors in [422, 718] showed that using LLMs as a paraphraser can easily evade several AI-text detection methods, even in the scenario when the original AI-text had been watermarked. These

findings sparked a heated debate about whether and how we can successfully design a reliable AI-text detector. While [718] theoretically quantifies the best detector’s performance with respect to the total variation distance between AI-text and human-text distributions and argues that AI-text is difficult to detect, another work [100] proves that it is possible to obtain a reliable AI-text detector unless the human-text distribution is exactly the same as the AI-text distribution, based on an information-theoretical analysis (i.e., the sample complexity of Chernoff information and likelihood-ratio-based detectors).

To improve the robustness of AI-text detection, we emphasize RADAR [327], a framework for training a robust AI-text detector using adversarial learning. The methodology of RADAR draws inspiration from adversarial machine learning techniques that train a high-quality generator by introducing a discriminator to form a two-player game, such as generative adversarial networks (GANs) [260]. In RADAR, we introduce a paraphraser and a detector as two players with opposite objectives. The paraphraser’s goal is to generate realistic content that can evade AI-text detection, while the detector’s goal is to enhance AI-text detectability. In RADAR’s implementation, both the paraphraser and the detector are parametrized by separate language models. During training, the paraphraser learns to rewrite the text from a training corpus (generated by a target LLM from a human-text corpus) with the aim of decreasing the likelihood of AI-text prediction by the detector, whereas the detector aims to enhance the detection performance by learning to compare human-text v.s. AI-text from the training data and the paraphraser’s output. These two players iteratively update their model parameters until their respective validation loss becomes stable. Specifically, the paraphraser treats the prediction of the detector as a reward and uses Proximal Policy Optimization (PPO) [738] for updates. The detector updates its parameters based on a logistic loss function evaluated on the human-text and AI-text corpora (including the texts generated by the paraphraser). In the evaluation phase, the trained detector is deployed to predict the likelihood of AI-written content for any input instance.

16.2 Background and Related Work

AI-Text Detection The research in AI-text detection can be divided into three approaches.

- (i) Statistical methods: some statistics such as entropy [438], n-gram frequency, and perplexity are used as a threshold to discern AI-text. A typical example is GLTR [252], which exploits entropy, probability, and probability rank for detection. Another example is DetectGPT [591], which assumes that the machine-generated text always lies in the negative curvature region of the log probability of the LLM of interest. Based on this hypothesis, DetectGPT perturbs the input text with a mask-filling language model, such as T5 [683].

Then, AI-text detection is performed by comparing the log probability of the text and its infilled variants.

- (ii) **Classification methods:** AI-text detection is formulated as a binary classification task, and a classifier is trained for a target language model [352, 705, 772, 973]. For example, OpenAI trains its AI-text classifier with a RoBERTa-based model [772]. The developers collected samples from the WebText dataset¹ and labeled them as human-generated. Then, for each target GPT-2 model, they collected the generated samples and labeled them as machine-generated. Finally, they fine-tuned the pretrained RoBERTa-based model [772] for AI-text classification. More recently, with the appearance of CharGPT, OpenAI tuned a GPT model called AI-Classifier¹ using data from several sources. The human-written text comes from three sources: a new Wikipedia dataset, the WebText dataset collected in 2019, and a set of human demonstrations collected as part of training InstructGPT [628]. To collect machine-generated text, for the Wikipedia and WebText datasets, they truncated the articles sampled from the original corpus and used 34 models to generate article completion, pairing each generated text with the original article. For the demonstrations, they used a model to generate responses for each prompt and paired them with the corresponding human demonstrations. This detector was only accessible via a web interface since its release in January 2023, and it has been taken down since July 2023.
- (iii) **Watermark methods:** post-hoc watermarking techniques, such as rule-based methods [77, 383, 814] and deep-learning-based methods [169, 840], can be applied to an LLM. At inference time, [409] proposed a soft watermarking scheme to embed a watermark in each word of the generated sentence by dividing the vocabulary into different lists and sampling the next token in a differentiated manner. However, many existing AI-text detectors are shown to be significantly weakened by paraphrasing in [718]. We also remark that current watermarking techniques won't be sufficient to detect AI-generated text for opensource LLMs, because bad actors can download the released model weights and opt out of the option of using watermarks for the generated text. Therefore, there is a strong demand for robust AI-text detectors, regardless of text watermarks.

Adversarial Learning for Natural Language Generation The success of GAN [260] in the computer vision domain has motivated many studies in natural language generation. However, since text generation is a sequential sampling process that occurs in a discrete vocabulary space, it is difficult to directly train a text generator using back-propagation in an end-to-end manner [161, 570, 944, 966]. There are two common approaches to tackle this problem. The first one is to replace the discrete sampling operation with continuous approximation techniques [161, 944],

¹ <https://huggingface.co/datasets/openwebtext>.

such as Gumbel-Softmax [359, 562]. The second one is to view text generation as a decision-making process and cast the generator as a policy [570, 910, 926, 966]. A typical example is SeqGAN [966]. During generation, SeqGAN considers the generated tokens as the state and the next token to be generated as the action, and it adopts Monte Carlo search to collect reward signals from the discriminator. Instead of using a classifier as the discriminator, the Diversity-Promoting GAN [926] uses a unidirectional LSTM as the discriminator and combines both word-level and sentence-level rewards into training. TextGAIL [910] proposed an imitation learning paradigm in which the rewards of the human-written text are regarded as a constant value. Then, both the rewards from human-text and AI-text are used to optimize the generator with PPO. These works all used warm-up training for the generator with maximum likelihood estimation (MLE) on the probability of the generated text sequence. On the other hand, [570] trained a language GAN from scratch. Our proposed RADAR differs from these works in that we focus on training a robust AI-text detector with a tunable paraphraser.

16.3 RADAR: Robust AI-Text Detection Using Adversarial Learning

An overview of RADAR [327] is illustrated in Fig. 16.1. The RADAR framework consists of three neural-network-based language models (LMs): the target LM $\mathcal{T}_\theta$, the detector $\mathcal{D}_\phi$ and the paraphraser $\mathcal{G}_\sigma$, parameterized with θ , ϕ and σ , respectively. We note that $\mathcal{T}_\theta$ is frozen (no updates on θ) in the entire process. We summarize RADAR into three key steps:

- **Step 1 (Data preparation):** Before training, we build $\mathcal{M}$, the corpus of AI-text, by applying document completion based on the prefix span of text in the human-text corpus $\mathcal{H}$ using $\mathcal{T}_\theta$.
- **Step 2 (Paraphraser update):** We collect AI-text samples x_m from $\mathcal{M}$ and use $\mathcal{G}_\phi$ to do paraphrasing on x_m to generate paraphrased AI-text x_p to form a corpus $\mathcal{P}$. Then, we use the reward of x_p returned by the detector $\mathcal{D}_\theta$ to update the paraphraser $\mathcal{G}_\phi$ using PPO.
- **Step 3 (Detector update):** We use the human-text samples x_h from $\mathcal{H}$, the original AI-text samples x_m from $\mathcal{M}$, and the paraphrased AI-text samples x_p from $\mathcal{P}$ in step 2 to update the detector $\mathcal{D}_\theta$ with a logistic loss function.
- **Step 4 (Performance Validation and Evaluation):** During training, we use the test set of WebText as the validation dataset to estimate RADAR’s performance. For evaluation, we use $\mathcal{T}_\theta$ to generate AI-text for the evaluation dataset and to calculate RADAR’s detection AUROC score.




Fig. 16.1 Overview of RADAR. An AI-text corpus is first generated from a target (frozen) language model from a human-text corpus. In RADAR, we introduce a paraphraser (a tunable language model) and a detector (a separate tunable language model). In the training stage, the detector aims to discern human-text v.s. AI-text, while the paraphraser aims to rewrite AI-text to evade detection. The model parameters of the paraphraser and the detector are updated in an adversarial learning manner as described in Sect. 16.3. In the evaluation stage, the trained detector is deployed to predict the likelihood of AI-generated content for any input instance

Step 2 to Step 3 can be repeated until there is no improvement in the AUROC evaluated on the validation dataset. The nature of rivalry in adversarial learning and the introduced competition helps the detector to learn to be robust in detecting both original and paraphrased AI-text.

16.3.1 Training Paraphraser via Clipped PPO with Entropy Penalty

In RADAR, the goal of the paraphraser G_σ is to paraphrase the input machine-generated text x_m . We model the generation of paraphrased text as a decision-making process, taking x_m as the state and the output text x_p as the action. In particular, we optimize G_σ using the reward feedback from the detector D_ϕ with PPO. The output of $D_\phi(x_p)$ is the predicted likelihood of x_p being Human-text.

The reward returned by x_p and the log probability of the text x_p are defined in Eq. (16.1):

$$R(x_p, \phi) = \mathcal{D}_\phi(x_p) \in [0, 1]; \quad \log P_{\mathcal{G}_\sigma}(x_p|x_m) = \sum_{i=1}^N \log P_{\mathcal{G}_\sigma}(x_p^i|x_m, x_p^{1:i-1}), \quad (16.1)$$

where x_p^i means the i -th token in the sentence x_p of length N and $x_p^{1:i-1}$ represents the first $i - 1$ tokens in x_p ($x_p^{1:0}$ means the default starting token).

RADAR uses Clipped PPO with Entropy Penalty (cppo-ep) in RADAR to optimize $\mathcal{G}_\sigma$. Let $\text{clip}(\cdot, a, b)$ denote a value-clipping operation with a lower limit a and an upper limit b , $r(\sigma, x_m, x_p)$ be the importance sampling ratio between a new policy $\mathcal{G}_\sigma$ and an old policy $\mathcal{G}_{\sigma'}$, and $(x_m, x_p) \sim P_{\mathcal{G}_{\sigma'}}$ be a state-action pair sampled from $\mathcal{G}_{\sigma'}$. The loss of cppo-ep is defined as:

$$L_{\mathcal{G}}(\sigma) = \mathbb{E}_{(x_m, x_p) \sim P_{\mathcal{G}_{\sigma'}}} \underbrace{-\min\{\text{clip}(r(\sigma, x_m, x_p), 1 - \epsilon, 1 + \epsilon), r(\sigma, x_m, x_p)\} \cdot A(x_p, \phi)}_{L_{\mathcal{A}}} \underbrace{-\gamma S(\sigma)}_{L_{\mathcal{E}}} \quad (16.2)$$

where $\mathbb{E}$ denotes expectation, ϵ is a parameter used in clipping to avoid the importance ratio r from being too large, $A(x_p, \phi)$ is the advantage item of the paraphrased text x_p obtained by applying normalization to $R(x_p, \phi)$ across the entire PPO sample buffer $\mathcal{B}$. $S(\sigma) = \mathbb{E}_{(x_m, x_p) \sim P_{\mathcal{G}_{\sigma'}}} -P_{\mathcal{G}_\sigma}(x_p|x_m) \log P_{\mathcal{G}_\sigma}(x_p|x_m)$, which is an entropy term introduced to encourage $\mathcal{G}_\sigma$ to explore more diverse generation policy. γ is a coefficient to control the ratio between $L_{\mathcal{A}}$ and $L_{\mathcal{E}}$, in order to make a balance between advantage ($L_{\mathcal{A}}$) and diversity ($L_{\mathcal{E}}$) when paraphrasing.

16.3.2 Training Detector via Reweighted Logistic Loss

In a typical GAN training process, the discriminator receives an equal amount of positive and negative samples in each step, assuring an in-batch sample balance. However, in RADAR, by construction, the number of AI-text samples is twice the number of human-text samples, because each x_h from the human-text corpus $\mathcal{H}$ is paired with a sample x_m from the original AI-text corpus $\mathcal{M}$ as well as a paraphrased sample x_p generated by the paraphraser $\mathcal{G}_\phi$. To handle this in-batch imbalance problem, we use a reweighted logistic loss function to optimize the detector D_ϕ , as

described in Eq. (16.3):

$$\begin{aligned}
 L_{\mathcal{D}}(\phi) = & \underbrace{-\mathbb{E}_{x_h \sim \mathcal{H}} \log \mathcal{D}_{\phi}(x_h)}_{L_{\mathcal{H}}: \text{loss on human-text}} + \underbrace{\lambda \mathbb{E}_{x_m \sim \mathcal{M}} - \log(1 - \mathcal{D}_{\phi}(x_m))}_{L_{\mathcal{M}}^1: \text{loss on original AI-text}} \\
 & + \underbrace{\lambda \mathbb{E}_{x_m \sim \mathcal{M}} - \log(1 - \mathcal{D}_{\phi}(\mathcal{G}_{\sigma}(x_m)))}_{L_{\mathcal{M}}^2: \text{loss on paraphrased AI-text}}
 \end{aligned} \tag{16.3}$$

Recall that $\mathcal{D}_{\phi}(x) \in [0, 1]$ is the predicted probability of an input instance x being Human-text. $L_{\mathcal{H}}$ is the loss to improve the correctness of predicting $x_h \sim \mathcal{H}$ as human-written. $L_{\mathcal{M}} = L_{\mathcal{M}}^1 + L_{\mathcal{M}}^2$, where $L_{\mathcal{M}}^1$ and $L_{\mathcal{M}}^2$ are used to avoid x_m and x_p from being predicted as human-text, respectively. λ is a coefficient ranging from 0 to 1. We introduce λ to adjust the proportion of AI-text components in the overall loss function to alleviate the effects of sample imbalance.

16.3.3 RADAR Algorithm

The entire training procedure of RADAR is summarized in Algorithm 10. For a given target LLM, RADAR returns a trained paraphraser and a trained detector through the designed training steps. In the evaluation phase, the detector is used to predict the likelihood of AI-text for any input instance.

16.4 Performance Evaluation

16.4.1 Experiment Setup

Datasets and Metrics For training, we sampled 160K documents from Web-Text [258] to build the human-text corpus $\mathcal{H}$. Then, we build the original AI-text corpus $\mathcal{M}$ from $\mathcal{H}$ using a target language model $\mathcal{T}_{\theta}$, which performs text completion using the first 30 tokens as the prompt and limits the sentence length to be 200 tokens. For evaluation, we select four human-text datasets covering different domains. Following [591], we use Xsum, SQuAD, and Reddit Writing-Prompts (WP) to test a detector’s ability to detect fake news, avoid academic fraud, and identify machine-generated literature innovation, respectively. In addition, we also use the non-native-authored TOEFL dataset (TOFEL) [493] to evaluate a detector’s bias when encountering non-native-authored English text. We report the area under the receiver operating characteristic curve (AUROC) score by varying the detector’s threshold as the performance measure (higher is better), which captures the relationship between the true positive rate and the false positive rate.

Algorithm 10 RADAR: robust AI-text detection via adversarial learning

```

1: Data initialization:
2: Collect human-written text to build human-text corpus  $\mathcal{H}$ 
3: Select a target language model  $\mathcal{T}_\theta$  to perform document completion on  $\mathcal{H}$  to build the
   corresponding AI-text corpus  $\mathcal{M}$ 
4: Build a replay buffer  $\mathcal{B}$  to store samples temporarily collected for training
5: Build a validation dataset  $\mathcal{V}$  from  $\mathcal{H}$  and  $\mathcal{M}$ 
6: Model initialization:
7: Detector  $\mathcal{D}_\phi \leftarrow \phi_{\text{pretrain}}$  (a pretrained language model)
8: Paraphraser  $\mathcal{G}_\sigma \leftarrow \sigma_{\text{pretrain}}$  (a pretrained language model)
9: Model training:
10: for  $i = 1$  : maximum step do
11:   Sample  $x_h$  and its corresponding  $x_m$  from  $\mathcal{H}$  and  $\mathcal{M}$  respectively
12:   Use  $\mathcal{G}_\sigma$  to paraphrase  $x_m$  and generate  $x_p$ 
13:   Collect reward  $R(x_p, \phi)$  as in Eq. (16.1)
14:   Normalize  $R(x_p, \phi)$  to compute the advantage function  $A(x_p, \phi)$  used in Eq. (16.2)
15:   Fill  $\mathcal{B}$  with  $(x_h, x_m, x_p, A(x_p, \phi))$ 
16:    $\sigma' \leftarrow \sigma$  # initialize the old policy  $\sigma'$  as the current policy  $\sigma$ 
17:   for  $(x_h, x_m, x_p, A(x_p, \phi)) \in \mathcal{B}$  do
18:     Compute the log probability  $\log P_{\mathcal{G}_\sigma}(x_p|x_m)$  and  $\log P_{\mathcal{G}'_\sigma}(x_p|x_m)$  using Eq. (16.1)
19:     Update  $\mathcal{G}_\sigma$  using Eq. (16.2)
20:   end for
21:   for  $(x_h, x_m, x_p, A(x_p, \phi)) \in \mathcal{B}$  do
22:     Update  $\mathcal{D}_\phi$  using Eq. 16.3
23:   end for
24:   Clear  $\mathcal{B}$ 
25:   Evaluate AUROC of  $\mathcal{D}_\phi$  on the validation dataset  $\mathcal{V}$ 
26: end for
27: Detector  $\mathcal{D}_\phi \leftarrow \phi_{\text{best}}$  (the detector model with the best AUROC on the validation dataset)
28: Paraphraser  $\mathcal{G}_\sigma \leftarrow \sigma_{\text{best}}$  (the paraphraser model which pairs with  $\phi_{\text{best}}$ )
29: Return  $\mathcal{D}_\phi$  and  $\mathcal{G}_\sigma$ 

```

Comparisons We compare RADAR with various detection methods. These methods include the OpenAI (RoBERTa) model which is fine-tuned on WebText [258] and GPT-2 [83] generations, as well as the statistical approaches including log probability, rank, log rank, entropy, and DetectGPT [252, 438, 591]. Specifically, we implemented DetectGPT using the trained T5-large model as the mask-filling model and performed 10 perturbations for each sentence to be detected.

Large Language Models For the target LLM $\mathcal{T}_\theta$, we select 4 pairs of LLMs and summarize them in Table 16.1. Each pair contains an open-source LLM and its fine-tuned version via instruction-tuning.

Table 16.1 Summary of the studied large language models

Parameter count	Model name	Organization	Pretrain data	Instruction fine-tune data
3B	Pythia-2.8B	EleutherAI	The pile ²	✗
	Dolly-V2-3B	Databricks		databricks-dolly-15k ³
5B	Palmyra-base	Writer	Writer’s custom dataset	✗
	Camel-5B	Writer		70K instruction-response records by Writer Linguist team
6B	GPT-J-6B	EleutherAI	The pile	✗
	Dolly-V1-6B	Databricks		Stanford Alpaca 52K instruction-following demonstrations ⁴
7B	LLaMA-7B	Meta	Various sources ⁵	✗
	Vicuna-7B	LMsys		70K conversations collected from ShareGPT ⁶

Paraphrase Configurations We consider two settings: *without (w/o) paraphrasing* and *with paraphrasing*. To prepare the machine-generated text for evaluation, for the w/o paraphrasing setting, we use the original AI-text corpus $\mathcal{M}$ generated by a target LLM based on an evaluation dataset. For the with paraphrasing setting, we define two types of paraphrasing: *seen paraphraser* and *unseen paraphraser*. The seen paraphraser refers to the paraphraser $\mathcal{G}_\sigma$ returned by RADAR. The unseen paraphraser means a new paraphraser that has not participated in training the detector of RADAR. We used the OpenAI API service of GPT-3.5-Turbo as the default unseen paraphraser. The prompt we used for paraphrasing is “Enhance word choices to make the sentence sound more like a human”, as inspired by Liang et al. [493].

Implementation Details We provide the detailed setups when implementing Algorithm 10. We build a PPO buffer $\mathcal{B}$ that can temporarily store 256 pairs of data for subsequent training. We use the pre-trained T5-large and RoBERTa-large models as the initialization of $\mathcal{G}_\sigma$ and $\mathcal{D}_\phi$ respectively. During training, we set the batch size to 32 and train the models until the validation loss converges. We use AdamW as the optimizer with the initial learning rate set to 1e-5 and use linear decay for both $\mathcal{G}_\sigma$ and $\mathcal{D}_\phi$. We set $\lambda = 0.5$ for sample balancing in Eq. 16.3 and set $\gamma = 0.01$ in Eq. 16.2. We follow the same construction principle of the training

² <https://huggingface.co/datasets/EleutherAI/pile>.

³ <https://huggingface.co/datasets/databricks/databricks-dolly-15k>.

⁴ https://github.com/tatsu-lab/stanford_alpaca/blame/main/alpaca_data.json.

⁵ Collected from CCNet [67%], C4 [15%], GitHub [4.5%], Wikipedia [4.5%], Books [4.5%], ArXiv [2.5%], Stack Exchange [2%].

⁶ <https://sharegpt.com/>.

Table 16.2 AUROC score averaged over 8 target LLMs. RADAR-seen paraphraser means the paraphraser used in RADAR ($\mathcal{G}_\sigma$). RADAR-unseen paraphraser is OpenAI’s GPT-3.5-Turbo API. The notations $\{\oplus, \ominus\}$ denote the best/second-best method for each dataset

Evaluation schema	Method	Evaluation dataset				
		Xsum	SQuAD	WP	TOFEL	Average
w/o Paraphraser	log p	0.882	0.868	0.967 $\ominus$	0.832	0.887
	rank	0.722	0.752	0.814	0.731	0.755
	log rank	0.902	0.893 $\ominus$	0.975 $\oplus$	0.847 $\ominus$	0.904 $\oplus$
	entropy	0.536	0.521	0.296	0.534	0.472
	DetectGPT	0.874	0.790	0.883	0.919 $\oplus$	0.867
	OpenAI (RoBERTa)	0.953 $\oplus$	0.914 $\oplus$	0.924	0.810	0.900 $\ominus$
	RADAR	0.934 $\ominus$	0.825	0.847	0.820	0.856
RADAR-seen paraphraser	log p	0.230	0.156	0.275	0.130	0.198
	rank	0.334	0.282	0.357	0.163	0.284
	log rank	0.245	0.175	0.281	0.134	0.209
	entropy	0.796	0.845 $\ominus$	0.763	0.876 $\ominus$	0.820 $\ominus$
	DetectGPT	0.191	0.105	0.117	0.177	0.159
	OpenAI (RoBERTa)	0.821 $\ominus$	0.842	0.892 $\ominus$	0.670	0.806
	RADAR	0.920 $\oplus$	0.927 $\oplus$	0.908 $\oplus$	0.932 $\oplus$	0.922 $\oplus$
RADAR-unseen paraphraser	log p	0.266	0.343	0.641	0.438	0.422
	rank	0.433	0.436	0.632	0.342	0.461
	log rank	0.282	0.371	0.632	0.421	0.426
	entropy	0.779	0.710 $\ominus$	0.499	0.618	0.651 $\ominus$
	DetectGPT	0.360	0.384	0.609	0.630 $\ominus$	0.434
	OpenAI (RoBERTa)	0.789 $\ominus$	0.629	0.726 $\ominus$	0.364	0.627
	RADAR	0.955 $\oplus$	0.861 $\oplus$	0.851 $\oplus$	0.763 $\oplus$	0.857 $\oplus$

dataset to create the 4 evaluation datasets based on Xsum, SQuAD, WP, and TOFEL. Experiments were run on 2 GPUS (NVIDIA Tesla V100 32 GB).

16.4.2 Performance Evaluation and Comparison with Existing Methods

We run three groups of experiments (w/o paraphraser, seen paraphraser, and unseen paraphraser) and report the overall results of RADAR and the compared methods on all 4 datasets in Table 16.2. The reported AUROC scores are averaged over the 8 considered LLMs. In the relatively easy case of without paraphrasing, most detectors attain good AUROC scores. RADAR attains a comparable performance (0.856) to the best existing detector (log rank, 0.904). The slightly worse performance of RADAR can be explained by the tradeoff in enhancing AI-text detection against paraphrasing.




Fig. 16.2 Performance evaluation (AUROC) of 8 LLMs over 4 human-text datasets. *w/o paraphraser* means the evaluation with the original AI-text corpora (the yellow bin $\mathcal{M}$ in Fig. 16.1). *RADAR-Unseen paraphraser* means the evaluation with the paraphrased AI-text (the green bin $\mathcal{P}$ in Fig. 16.1) generated from an independent paraphraser (OpenAI’s GPT-3.5-Turbo API) that is not used in RADAR. The black error bar represents the standard deviation of the detection AUROCs across 8 LLMs

When facing paraphrasing, all existing methods except entropy show significant performance degradation. The drop in AUROC compared to the *w/o* paraphrasing case ranges from 10.4 to 81.7%. While entropy is shown to be more robust to paraphrasing, its AUROC score can be quite low. On the contrary, RADAR demonstrates robust and superior detection power, attaining the best performance on every dataset. As shown in Fig. 16.2, the average AUROC score of RADAR (0.857) improves the best existing method (entropy, 0.651) by 31.64% against the unseen paraphraser. On average, RADAR is more robust to the seen paraphraser than the unseen paraphraser, because the seen paraphraser is what is used to train the detector in RADAR. More importantly, the detection performance of RADAR is stable across different paraphrasing schema, suggesting that RADAR can successfully mitigate the performance drop in AI-text detection.

16.4.3 AI-Text Detection Transferability of RADAR

We explore the AI-text detection transferability of RADAR between the 8 LLMs and report the ratio $F(A,B)=\text{AUROC}(A,B)/\text{AUROC}(B,B)$ for each LLM pair (A,B) , where $\text{AUROC}(A,B)$ means using the RADAR’s detector trained on model A to evaluate the AI-text generated by model B. A larger ratio means better transferability from A to B. Figure 16.3 shows the matrix of pairwise detection transferability and the bar chart of the holistic detection transferability to all the 8 LLMs in the without and unseen paraphrasing settings. We highlight two key observations as follows.




Fig. 16.3 RADAR’s detection transferability between pairs of 8 LLMs in Table 16.1. In the matrix, each row is the source LLM (model A) for training the detector, and each column is the target LLM (model B) for evaluation. The reported value in the matrix represents the detection transferability from A to B. A larger value indicates better transferability. The bar chart shows the row-wise sum of the matrix, indicating the holistic transferability of each source LLM. (a) w/o paraphraser. (b) GPT-3.5-Turbo paraphraser

(I) Instruction-Tuned Models Have Better Detection Transferability Partitioning the LLMs into two groups, we find that the detector targeting an instruction-tuned LLM (top 4 rows) generally transfers better than the detector targeting the corresponding LLM without instruction-tuning (bottom 4 rows). Take the pair (Vicuna-7B, LLaMA-7B) as an example, we can see that without paraphrasing, $F(\text{Vicuna-7B}, \text{LLaMA})$ can reach up to 95.0%. On the other hand, $F(\text{LLaMA-7B}, \text{Vicuna-7B})$ can only account for 68.2%. Sorting the detectors according to the holistic detection transferability (which is presented in the bar chart), we can see the top-3 detectors are all trained with the instruction-tuned LLMs. A similar conclusion can be made for the with paraphrasing setting. Moreover, there is no obvious trend between the target LLM size and the resulting detection performance. The effect of instruction tuning on transferability is more prominent than model size.

(II) RADAR Achieves Better Detection Transferability Against Paraphrasing Another interesting finding is that RADAR’s transferability is generally improved when paraphrasing is in place. Comparing the two bar charts in Fig. 16.3a, b, the average holistic detection transferability (over all LLMs) is increased by 11.6%. Except for LLaMA-7B (3.8% drop) and GPT-J-6B (1.4% drop), all other LLMs’ holistic transferability scores are improved from 2.4% (Palmyra-base) to 47.6% (Camel-5B).

Chapter 17

Backdoor Risks in Diffusion Models



Abstract Diffusion Models (DMs) are state-of-the-art generative models that learn a reversible corruption process from iterative noise addition and denoising. They are the backbone of many generative AI applications, such as text-to-image conditional generation. The chapter explores the backdoor risks of DMs, which can be viewed as a type of output manipulation attack triggered by a maliciously embedded pattern at model input.

17.1 Introduction

In recent years, diffusion models (DMs) [47, 184, 299–301, 387, 520, 548, 549, 707, 769, 774, 776–778, 987] trained with large-scale datasets [735, 736] have emerged as a cutting-edge content generation AI tool, including image [184, 299, 301, 610, 687, 719], audio [418], video [303, 579], text [485], and text-to-speech [338, 365, 397, 659] generation. Even more, DMs are increasingly used in safety-critical tasks and content curation, such as reinforcement learning, object detection, and inpainting [48, 115, 127, 140, 361, 645, 876].

This chapter explores the risk of backdoor attacks on DMs. Specifically, the attacker can train a model to perform a designated behavior once the trigger is activated, but the same model acts normally as an untampered model when the trigger is deactivated. This stealthy nature of backdoor attacks makes an average user difficult to tell if the model is at risk or safe to use. The implications of such backdoor injection attacks include content manipulation (e.g. generating inappropriate content for image inpainting), falsification (e.g. spoofing attacks), and model watermarking (by viewing the embedded trigger as a watermark query). Further, the attacker can also use backdoored DMs to generate biased or adversarial datasets at scale [132, 189], which may indirectly cause future models to become problematic and unfair.

This chapter presents **VillanDiffusion**, a unified backdoor attack framework for DMs in [146]. It covers (1) generalization to both denoising diffusion models like DDPM [299, 769] and score-based models like NCSN [776–778]; (2) extension to various advanced training-free samplers like DPM Solver [548, 549], PNDM




Fig. 17.1 Overview of VillanDiffusion, a unified backdoor attack framework for DMs proposed in [146]. An image or a prompt with a designed trigger pattern can trigger a backdoored DM to generate a specific target image (here, the cat image)

[520], UniPC [1016] and DEIS [987] without modifying the samplers; and (3) demonstration that a text-to-image DM can be backdoored in the prompt space even if the text encoder is untouched.

As illustrated in Fig. 17.1, we categorize the DMs based on three perspectives: (1) schedulers, (2) samplers, and (3) conditional and unconditional generation. We summarize the key features of the unified framework as follows.

- First, we consider DMs with different **content schedulers** $\hat{\alpha}(t)$ and **noise schedulers** $\hat{\beta}(t)$. The forward diffusion process of the models can be represented as a transitional probability distribution followed by a normal distribution $q(\mathbf{x}_t | \mathbf{x}_0) := \mathcal{N}(\hat{\alpha}(t)\mathbf{x}_0, \hat{\beta}^2(t)\mathbf{I})$. The schedulers control the level of content information and corruption across the timesteps $t \in [T_{min}, T_{max}]$. We also denote $q(\mathbf{x}_0)$ as the data distribution. To show the generalizability of our framework, we discuss two major branches of DMs: DDPM [299] and Score-Based Models [776–778]. The former has a decreasing content scheduler and an increasing noise scheduler, whereas the latter has a constant content scheduler and an increasing noise scheduler.
- Secondly, VillanDiffusion also considers different kinds of samplers. In [548, 778], the generative process of DMs can be described as a reversed-time stochastic differential equation (SDE):

$$d\mathbf{x}_t = [\mathbf{f}(\mathbf{x}_t, t) - g^2(t)\nabla_{\mathbf{x}_t} \log q(\mathbf{x}_t)]dt + g(t)d\bar{\mathbf{w}} \quad (17.1)$$

The reverse-time SDE can also be written as a reverse-time ordinary differential equation (ODE) in (17.2) with the same marginal probability $q(\mathbf{x}_t)$. We found that the additional coefficient $\frac{1}{2}$ will cause BadDiffusion [145], a prior art, to fail on the ODE samplers, including DPM-Solver [548] and DDIM [774].

$$d\mathbf{x}_t = [\mathbf{f}(\mathbf{x}_t, t) - \frac{1}{2}g^2(t)\nabla_{\mathbf{x}_t} \log q(\mathbf{x}_t)]dt \quad (17.2)$$

- Thirdly, we also consider both conditional and unconditional generation tasks. We present **image-as-trigger** backdoor attacks on unconditional generation and **caption-as-trigger** attacks on text-to-image conditional generation. Compared to [145], which only studies one DM (DDPM) on unconditional generation with image triggers, VillanDiffusion can generalize to various DMs, including DDPM [299] and the score-based models [776–778]. In [132], only DDPM and DDIM [774] are studied and the attackers are allowed to modify the samplers. VillanDiffusion covers a diverse set of off-the-self samplers without assuming the attacker has control over the samplers.

17.2 Background and Related Work

Diffusion Models DMs are designed to learn the reversed diffusion process which is derived from a tractable forward corruption process [769, 778]. Since the diffusion process is well-studied and reversible, it does not require special architecture design like flow-based models [187, 405, 702]. Generally, most diffusion models follow different schedulers to determine the Gaussian noise and the content levels at different timesteps. Commonly used diffusion models are DDPM [299], score-based models [776, 777], and VDM [406], etc.

Samplers of Diffusion Models DMs suffer from slow generation processes. Recent works mainly focus on sampling acceleration like PNDM [520] and EDM [387], which treat the diffusion process as an ODE and apply high-order approximation to reduce the error. Moreover, samplers including UniPC [1016], DEIS [987], DPM Solver [548], and DPM-Solver++ [549] leverage the semi-linear property of diffusion processes to derive a more precise approximation. On the other hand, DDIM [774] discards Markovian assumption to accelerate the generative process. Another training-based method is distilling DMs, such as [724].

Backdoor Attack on Diffusion Models BadDiffusion [145] backdoors DDPM with an additional correction term on the mean of the forward diffusion process without any modification on the samplers. TrojDiff [132] assumes the attacker can access both training procedures and samplers and apply correction terms on DDPM [299] and DDIM [774] to launch the attack. The work [784] backdoors text-to-image DMs via altering the text encoder instead of the DMs.

Threat Model and Attack Scenario With ever-increasing training costs in scale and model size, adopting pre-trained models become a common choice for most users and developers. We follow [145] to formulate the attack scenario with two parties: (1) an *attacker*, who releases the backdoored models on the web, and (2) a *user*, who downloads the pre-trained models from third-party websites like HuggingFace. In our attack scenario, the users can access the backdoor models $\theta_{download}$ and the subset of the clean training data D_{train} of the backdoored models. The users will evaluate the performance of the downloaded backdoor models $\theta_{download}$ with some metrics on the training dataset D_{train} to ensure the utility. For image generative models, the FID [296] and IS [723] scores are widely used metrics. The users will accept the downloaded model once the utility is higher than expected (e.g. the utility of a clean model). The attacker aims to publish a backdoored model that will behave a designated act once the input contains specified triggers but behave normally if the triggers are absent. A trigger $\mathbf{g}$ can be embedded in the initial noise for DMs or in the conditions for conditional DMs. The designated behavior is to generate a target image $\mathbf{y}$. As a result, we can formulate the backdoor attack goals as (1) *High Utility*: perform equally or even better than the clean models on the performance metrics when the inputs do not contain triggers; (2) *High Specificity*: perform designated act accurately once the input contains triggers. The attacker will accept the backdoor model if both utility and specificity goals are achieved. For image generation, we use the FID [296] score to measure the utility and use the mean squared error (MSE) to quantify the specificity.

17.3 VillanDiffusion: A Unified Backdoor Attack Framework

17.3.1 Backdoor Unconditional Diffusion Models as a Distribution Mapping Problem

Clean Forward Diffusion Process Generative models aim to generate data that follow a ground-truth data distribution $q(\mathbf{x}_0)$ from a simple prior distribution π . Thus, we can treat it as a distribution mapping from the prior distribution π to the data distribution $q(\mathbf{x}_0)$. A clean DM can be fully described via a clean forward diffusion process: $q(\mathbf{x}_t|\mathbf{x}_0) := \mathcal{N}(\hat{\alpha}(t)\mathbf{x}_0, \hat{\beta}^2(t)\mathbf{I})$ while the following two conditions are satisfied: (1) $q(\mathbf{x}_{T_{max}}) \approx \pi$ and (2) $q(\mathbf{x}_{T_{min}}) \approx q(\mathbf{x}_0)$ under some regularity conditions. Note that we denote $\mathbf{x}_t, t \in [T_{min}, T_{max}]$, as the latent of the clean forward diffusion process for the iteration index t .

Backdoor Forward Diffusion Process with Image Triggers When backdooring unconditional DMs, we use a chosen pattern as the trigger $\mathbf{g}$. Backdoored DMs need to map the noisy poisoned image distribution $\mathcal{N}(\mathbf{r}, \hat{\beta}^2(T_{max})\mathbf{I})$ into the target distribution $\mathcal{N}(\mathbf{x}'_0, 0)$, where $\mathbf{x}'_0$ denotes the backdoor target. Thus, a backdoored DM can be described as a backdoor forward diffusion process $q(\mathbf{x}'_t|\mathbf{x}'_0) := \mathcal{N}(\hat{\alpha}(t)\mathbf{x}'_0 +$

$\hat{\rho}(t)\mathbf{r}, \hat{\beta}^2(t)\mathbf{I})$ with two conditions: (1) $q(\mathbf{x}'_{T_{max}}) \approx \mathcal{N}(\mathbf{r}, \hat{\beta}^2(T_{max})\mathbf{I})$ and (2) $q(\mathbf{x}'_{T_{min}}) \approx \mathcal{N}(\mathbf{x}'_0, 0)$. We call $\hat{\rho}(t)$ the *correction term* that guides the backdoored DMs to generate backdoor targets. Note that we denote the latent of the backdoor forward diffusion process as $\mathbf{x}'_t, t \in [T_{min}, T_{max}]$, backdoor target as $\mathbf{x}'_0$, and poison image as $\mathbf{r} := \mathbf{M} \odot \mathbf{g} + (1 - \mathbf{M}) \odot \mathbf{x}$, where $\mathbf{x}$ is a clean image sampled from the clean data $q(\mathbf{x}_0)$, $\mathbf{M} \in \{0, 1\}$ is a binary mask indicating, the trigger g is stamped on $\mathbf{x}$, and $\odot$ means element-wise product.

Optimization Objective of the Backdoor Attack on Diffusion Models Consider the two goals of backdooring unconditional generative models: high utility and high specificity, we can achieve these goals by optimizing the marginal probability $p_\theta(\mathbf{x}_0)$ and $p_\theta(\mathbf{x}'_0)$ with trainable parameters θ . We formulate the optimization of the negative-log likelihood (NLL) objective in (17.3), where η_c and η_p denote the weight of utility and specificity goals, respectively.

$$\arg \min_{\theta} -(\eta_c \log p_\theta(\mathbf{x}_0) + \eta_p \log p_\theta(\mathbf{x}'_0)) \quad (17.3)$$

17.3.2 Generalization to Various Schedulers

We expand on the optimization problem formulated in (17.3) with variational lower bound (VLBO) and provide a more general computational scheme. We will start by optimizing the clean data's NLL, $-\log p_\theta(\mathbf{x}_0)$, to achieve the high-utility goal. Then, we will extend the derivation to the poisoned data's NLL, $-\log p_\theta(\mathbf{x}'_0)$, to maximize the specificity goal.

The Clean Reversed Transitional Probability Assume the data distribution $q(\mathbf{x}_0)$ follows the empirical distribution. From the variational perspective, minimizing the VLBO in (17.4) of a DM with trainable parameters θ is equivalent to reducing the NLL in (17.3). Namely,

$$\begin{aligned} -\log p_\theta(\mathbf{x}_0) &= -\mathbb{E}_q[\log p_\theta(\mathbf{x}_0)] \\ &\leq \mathbb{E}_q\left[\mathcal{L}_T(\mathbf{x}_T, \mathbf{x}_0) + \sum_{t=2}^T \mathcal{L}_t(\mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_0) - \mathcal{L}_0(\mathbf{x}_1, \mathbf{x}_0)\right] \end{aligned} \quad (17.4)$$

Denote $\mathcal{L}_t(\mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_0) = D_{\text{KL}}(q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) \parallel p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t))$, $\mathcal{L}_T(\mathbf{x}_T, \mathbf{x}_0) = D_{\text{KL}}(q(\mathbf{x}_T|\mathbf{x}_0) \parallel p_\theta(\mathbf{x}_T))$, and $\mathcal{L}_0(\mathbf{x}_1, \mathbf{x}_0) = \log p_\theta(\mathbf{x}_0|\mathbf{x}_1)$, where $D_{\text{KL}}(q||p) = \int_{\mathbf{x}} q(\mathbf{x}) \log \frac{q(\mathbf{x})}{p(\mathbf{x})}$ is the KL-Divergence. Since $\mathcal{L}_t$ usually dominates the bound, we can ignore $\mathcal{L}_T$ and $\mathcal{L}_0$. Because the ground-truth reverse transitional probability $q(\mathbf{x}_{t-1}|\mathbf{x}_t)$ is intractable, to compute $\mathcal{L}_t$, we can use a tractable conditional reverse transition $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$ to approximate it with a simple equation $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) = q(\mathbf{x}_t|\mathbf{x}_{t-1}) \frac{q(\mathbf{x}_{t-1}|\mathbf{x}_0)}{q(\mathbf{x}_t|\mathbf{x}_0)}$ based on the Bayesian and the Markovian rule. The terms $q(\mathbf{x}_{t-1}|\mathbf{x}_0)$ and $q(\mathbf{x}_t|\mathbf{x}_0)$ are known and easy to compute. To compute $q(\mathbf{x}_t|\mathbf{x}_{t-1})$

in close form, DDPM [299] proposes a well-designed scheduler. However, it does not apply to other scheduler choices like score-based models [776–778]. Consider the generalizability, we use numerical methods to compute the forward transition $q(\mathbf{x}_t | \mathbf{x}_{t-1}) := \mathcal{N}(k_t \mathbf{x}_{t-1}, w_t^2 \mathbf{I})$ since the forward diffusion process follows Gaussian distribution. Then, we reparametrize $\mathbf{x}_t$ based on the recursive definition: $\bar{\mathbf{x}}_t(\mathbf{x}, \epsilon_t) = k_t \bar{\mathbf{x}}_{t-1}(\mathbf{x}, \epsilon_{t-1}) + w_t \epsilon_t$ as described in (17.5).

$$\begin{aligned} \bar{\mathbf{x}}_t(\mathbf{x}_0, \epsilon_t) &= k_t \bar{\mathbf{x}}_{t-1}(\mathbf{x}_0, \epsilon_{t-1}) + w_t \epsilon_t = k_t(k_{t-1} \bar{\mathbf{x}}_{t-2}(\mathbf{x}_0, \epsilon_{t-2}) + w_{t-1} \epsilon_{t-1}) + w_t; \\ \epsilon_t &= \prod_{i=1}^t k_i \mathbf{x}_0 + \sqrt{\sum_{i=1}^{t-1} \left(\left(\prod_{j=i+1}^t k_j \right) w_i \right)^2} + w_t^2 \cdot \epsilon, \quad \forall t, \epsilon, \epsilon_t \stackrel{i.i.d}{\sim} \mathcal{N}(0, \mathbf{I}) \end{aligned} \quad (17.5)$$

Recall the reparametrization of the forward diffusion process: $\mathbf{x}_t(\mathbf{x}_0, \epsilon) = \hat{\alpha}(t) \mathbf{x}_0 + \hat{\beta}(t) \epsilon$, we can derive $\hat{\alpha}(t) = \prod_{i=1}^t k_i$ and $\hat{\beta}(t) = \sqrt{\sum_{i=1}^{t-1} \left(\left(\prod_{j=i+1}^t k_j \right) w_i \right)^2} + w_t^2$. Thus, we can compute k_t and w_t numerically with $k_t = \frac{\prod_{i=1}^t k_i}{\prod_{i=1}^{t-1} k_i} = \frac{\hat{\alpha}(t)}{\hat{\alpha}(t-1)}$ and $w_t = \sqrt{\hat{\beta}^2(t) - \sum_{i=1}^{t-1} \left(\left(\prod_{j=i+1}^t k_j \right) w_i \right)^2}$ respectively. With the numerical solutions k_t and w_t , we can follow the similar derivation of DDPM [299] and compute the conditional reverse transition in (17.6) with $a(t) = \frac{k_t \hat{\beta}^2(t-1)}{k_t^2 \hat{\beta}^2(t-1) + w_t^2}$ and $b(t) = \frac{\hat{\alpha}(t-1) w_t^2}{k_t^2 \hat{\beta}^2(t-1) + w_t^2}$:

$$q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0) := \mathcal{N}(a(t) \mathbf{x}_t + b(t) \mathbf{x}_0, s^2(t) \mathbf{I}), \quad s(t) = \sqrt{\frac{b(t)}{\hat{\alpha}(t)}} \hat{\beta}(t) \quad (17.6)$$

Finally, based on (17.6), we can follow the derivation of DDPM [299] and derive the denoising loss function in (17.7) to maximize the utility. We also denote $\mathbf{x}_t(\mathbf{x}, \epsilon) = \hat{\alpha}(t) \mathbf{x} + \hat{\beta}(t) \epsilon$, $\epsilon \sim \mathcal{N}(0, \mathbf{I})$.

$$L_c(\mathbf{x}, t, \epsilon) := \|\epsilon - \epsilon_\theta(\mathbf{x}_t(\mathbf{x}, \epsilon), t)\|^2 \quad (17.7)$$

On the other hand, we can also interpret (17.7) as a denoising score matching loss, which means the expectation of (17.7) is proportional to the score function, i.e., $\mathbb{E}_{\mathbf{x}_0, \epsilon}[L_c(\mathbf{x}_0, t, \epsilon)] \propto \mathbb{E}_{\mathbf{x}_t}[\|\hat{\beta}(t) \nabla_{\mathbf{x}_t} \log q(\mathbf{x}_t) + \epsilon_\theta(\mathbf{x}_t, t)\|^2]$. We further derive the backdoor reverse transition as follows.

The Backdoor Reversed Transitional Probability. Following similar ideas, we optimize VLBO instead of the backdoor data's NLL in (17.8) as

$$\begin{aligned} -\log p_\theta(\mathbf{x}'_0) &= -\mathbb{E}_q[\log p_\theta(\mathbf{x}'_0)] \\ &\leq \mathbb{E}_q[\mathcal{L}_T(\mathbf{x}'_T, \mathbf{x}'_0) + \sum_{t=2}^T \mathcal{L}_t(\mathbf{x}'_t, \mathbf{x}'_{t-1}, \mathbf{x}'_0) - \mathcal{L}_0(\mathbf{x}'_1, \mathbf{x}'_0)] \end{aligned} \quad (17.8)$$

Denote the backdoor forward transition $q(\mathbf{x}'_t|\mathbf{x}'_{t-1}) := \mathcal{N}(k_t\mathbf{x}'_{t-1} + h_t\mathbf{r}, w_t^2\mathbf{I})$. With a similar parametrization trick, we can compute h_t as $h_t = \hat{\rho}(t) - \sum_{i=1}^{t-1} \left(\left(\prod_{j=i+1}^t k_j \right) h_i \right)$. Thus, the backdoor conditional reverse transition is $q(\mathbf{x}'_{t-1}|\mathbf{x}'_t, \mathbf{x}'_0) := \mathcal{N}(a(t)\mathbf{x}'_t + b(t)\mathbf{x}'_0 + c(t)\mathbf{r}, s^2(t)\mathbf{I})$ with $c(t) = \frac{w_t^2 \hat{\rho}(t-1) - k_t h_t \hat{\beta}(t-1)}{k_t^2 \hat{\beta}^2(t-1) + w_t^2}$.

17.3.3 Generalization to ODE and SDE Samplers

In Sect. 17.3.2, we have derived a general form for both clean and backdoor reversed transitional probability $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$ and $q(\mathbf{x}'_{t-1}|\mathbf{x}'_t, \mathbf{x}'_0)$. Since DDPM uses $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$ to approximate the intractable term $q(\mathbf{x}_{t-1}|\mathbf{x}_t)$, as we minimize the KL-divergence between the two reversed transitional probabilities $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$ and $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$ in $L_t(\mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_0)$, it actually forces the model with parameters θ to learn the joint probability $q(\mathbf{x}_{0:T})$, which is the discrete trajectory of a stochastic process. As a result, we can convert the transitional probability into a stochastic differential equation and interpret the optimization process as a score-matching problem [775]. With the Fokker-Planck [548, 778], we can describe the SDE as a PDE by differentiating the marginal probability on the timestep t . We can further generalize our backdoor attack to various ODE samplers in a unified manner, including DPM-Solver [548, 549], DEIS [987], PNDM [520], etc.

Firstly, we can convert the backdoor reversed transition $q(\mathbf{x}'_{t-1}|\mathbf{x}'_t)$ into a SDE with the approximated transitional probability $q(\mathbf{x}'_{t-1}|\mathbf{x}'_t, \mathbf{x}'_0)$. With reparametrization, $\mathbf{x}'_{t-1} = a(t)\mathbf{x}'_t + c(t)\mathbf{r} + b(t)\mathbf{x}'_0 + s(t)\epsilon$ in (17.3.2) and $\mathbf{x}'_t = \hat{\alpha}(t)\mathbf{x}'_0 + \hat{\rho}(t)\mathbf{r} + \hat{\beta}(t)\epsilon_t$ in (17.3.1), we can present the backdoor reversed process $q(\mathbf{x}'_{t-1}|\mathbf{x}'_t)$ as a SDE with $F(t) = a(t) + \frac{b(t)}{\hat{\alpha}(t)} - 1$ and $H(t) = c(t) - \frac{b(t)\hat{\rho}(t)}{\hat{\alpha}(t)}$:

$$d\mathbf{x}'_t = [F(t)\mathbf{x}'_t - \underbrace{G^2(t)(-\hat{\beta}(t)\nabla_{\mathbf{x}'_t} \log q(\mathbf{x}'_t) - \frac{H(t)}{G^2(t)}\mathbf{r})}_{\text{Backdoor Score Function}}]dt + G(t)\sqrt{\hat{\beta}(t)}d\bar{\mathbf{w}}, \quad (17.9)$$

where $G(t) = \sqrt{\frac{b(t)\hat{\beta}(t)}{\hat{\alpha}(t)}}$.

To describe the backdoor reversed SDE in a process with arbitrary stochasticity, based on the Fokker-Planck equation we further convert the SDE in (17.9) into another SDE in (17.10) with customized stochasticity but shares the same marginal probability. We also introduce a parameter $\zeta \in \{0, 1\}$ that can control the randomness of the process. ζ can also be determined by the samplers directly. The process (17.10) will reduce to an ODE when $\zeta = 0$. It will be an SDE when $\zeta = 1$.

$$d\mathbf{x}'_t = [F(t)\mathbf{x}'_t - \frac{1+\zeta}{2}G^2(t) \underbrace{(-\hat{\beta}(t)\nabla_{\mathbf{x}'_t} \log q(\mathbf{x}'_t) - \frac{2H(t)}{(1+\zeta)G^2(t)}\mathbf{r})}_{\text{Backdoor Score Function}}]dt + G(t)\sqrt{\zeta\hat{\beta}(t)}d\tilde{\mathbf{w}} \quad (17.10)$$

When we compare it to the learned reversed process of SDE in (17.11), we can see that the diffusion model ϵ_θ should learn the backdoor score function to generate the backdoor target distribution $q(\mathbf{x}'_0)$.

$$d\mathbf{x}_t = [F(t)\mathbf{x}_t - \frac{1+\zeta}{2}G^2(t)\epsilon_\theta(\mathbf{x}_t, t)]dt + G(t)\sqrt{\zeta\hat{\beta}(t)}d\tilde{\mathbf{w}} \quad (17.11)$$

As a result, the backdoor score function will be the learning objective of the DM with ϵ_θ . We note that one can further extend this framework to DDIM [774] and EDM [387], which have an additional hyperparameter to control the stochasticity of the generative process.

17.3.4 Unified Loss Function for Unconditional Generation with Image Triggers

Following the aforementioned analysis, to achieve the high-specificity goal, we can formulate the loss function as $\mathbb{E}_{\mathbf{x}_0, \mathbf{x}'_t} [||(-\hat{\beta}(t)\nabla_{\mathbf{x}'_t} \log q(\mathbf{x}'_t) - \frac{2H(t)}{(1+\zeta)G^2(t)}\mathbf{r}) - \epsilon_\theta(\mathbf{x}'_t, t)||^2] \propto \mathbb{E}_{\mathbf{x}_0, \mathbf{x}'_0, \epsilon} [||\epsilon - \frac{2H(t)}{(1+\zeta)G^2(t)}\mathbf{r} - \epsilon_\theta(\mathbf{x}'_t(\mathbf{x}'_0, \mathbf{r}, t), \epsilon)||^2]$ with reparametrization $\mathbf{x}'_t(\mathbf{x}, \mathbf{r}, \epsilon) = \hat{\alpha}(t)\mathbf{x} + \hat{\rho}(t)\mathbf{r} + \hat{\beta}(t)\epsilon$. Therefore, we can define the backdoor loss function as $L_p(\mathbf{x}, t, \epsilon, \mathbf{g}, \mathbf{y}, \zeta) := ||\epsilon - \frac{2H(t)}{(1+\zeta)G^2(t)}\mathbf{r}(\mathbf{x}, \mathbf{g}) - \epsilon_\theta(\mathbf{x}'_t(\mathbf{y}, \mathbf{r}(\mathbf{x}, \mathbf{g}), \epsilon), t)||^2$ where the parameter ζ will be 0 when backdooring ODE samplers and 1 when backdooring SDE samplers. Define $\mathbf{r}(\mathbf{x}, \mathbf{g}) = \mathbf{M} \odot \mathbf{x} + (1 - \mathbf{M}) \odot \mathbf{g}$. We derive the unified loss function for unconditional DMs in (17.12). We can also show that BadDiffusion [145] is just a special case of it with proper settings.

$$L_\theta^I(\eta_c, \eta_p, \mathbf{x}, t, \epsilon, \mathbf{g}, \mathbf{y}, \zeta) := \eta_c L_c(\mathbf{x}, t, \epsilon) + \eta_p L_p(\mathbf{x}, t, \epsilon, \mathbf{g}, \mathbf{y}, \zeta) \quad (17.12)$$

Algorithm 11 Backdoor unconditional DMs with image trigger

Inputs: Backdoor Image Trigger $\mathbf{g}$, Backdoor Target $\mathbf{y}$, Training dataset D , Training parameters θ , Sampler Randomness ζ

while not converge **do**

$\{\mathbf{x}, \eta_c, \eta_p\} \sim D$

$t \sim \text{Uniform}(\{1, \dots, T\})$

$\epsilon \sim \mathcal{N}(0, \mathbf{I})$

 Use gradient descent $\nabla_{\theta} L_{\theta}^I(\eta_c, \eta_p, \mathbf{x}, t, \epsilon, \mathbf{g}, \mathbf{y}, \zeta)$ to update θ

end while

We summarize the training algorithm in Algorithm 11. Note that every data point $\mathbf{e}^i = \{\mathbf{x}^i, \eta_c^i, \eta_p^i\}$, $\mathbf{e}^i \in D$ in the training dataset D consists of three elements: (1) clean training image $\mathbf{x}^i$, (2) clean loss weight η_c^i , and (3) backdoor loss weight η_p^i . The poison rate defined in BadDiffusion [145] can be interpreted as $\frac{\sum_{i=1}^N \eta_p^i}{|D|}$, where $\eta_p^i, \eta_c^i \in \{0, 1\}$. We also denote the training dataset size as $|D| = N$.

17.3.5 Generalization to Conditional Generation

To backdoor a conditional generative DM, we can optimize the joint probability $q(\mathbf{x}_0, \mathbf{c})$ with a condition $\mathbf{c}$ instead of the marginal $q(\mathbf{x}_0)$. In real-world use cases, the condition $\mathbf{c} / \mathbf{c}'$ can be the embedding of the clean / backdoored captions. The resulting generalized objective function becomes

$$\arg \min_{\theta} -(\eta_c \log p_{\theta}(\mathbf{x}_0, \mathbf{c}) + \eta_p \log p_{\theta}(\mathbf{x}'_0, \mathbf{c}')) \quad (17.13)$$

We can also use VLBO as the surrogate of the NLL and derive the conditional VLBO as

$$-\log p_{\theta}(\mathbf{x}_0, \mathbf{c}) \leq \mathbb{E}_q \left[\mathcal{L}_T^C(\mathbf{x}_T, \mathbf{x}_0, \mathbf{c}) + \sum_{t=2}^T \mathcal{L}_t^C(\mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_0, \mathbf{c}) - \mathcal{L}_0^C(\mathbf{x}_1, \mathbf{x}_0, \mathbf{c}) \right] \quad (17.14)$$

Denote $\mathcal{L}_T^C(\mathbf{x}_T, \mathbf{x}_0, \mathbf{c}) = D_{\text{KL}}(q(\mathbf{x}_T | \mathbf{x}_0) \parallel p_{\theta}(\mathbf{x}_T, \mathbf{c}))$, $\mathcal{L}_0^C(\mathbf{x}_1, \mathbf{x}_0, \mathbf{c}) = \log p_{\theta}(\mathbf{x}_0 | \mathbf{x}_1, \mathbf{c})$, and $\mathcal{L}_t^C(\mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_0, \mathbf{c}) = D_{\text{KL}}(q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0) \parallel p_{\theta}(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{c}))$. To compute $\mathcal{L}_t^C(\mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_0, \mathbf{c})$, we need to compute $q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0, \mathbf{c})$ and $p_{\theta}(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{c})$ first. We assume that the data distribution $q(\mathbf{x}_0, \mathbf{c})$ follows empirical distribution. Thus, using the same derivation as in Sect. 17.3.2, we can obtain the clean data's loss function $L_c^C(\mathbf{x}, t, \epsilon, \mathbf{c}) := \|\epsilon - \epsilon_{\theta}(\mathbf{x}_t(\mathbf{x}, \epsilon), t, \mathbf{c})\|^2$ and we can derive the caption-trigger backdoor loss function as

$$L_{\theta}^{CC}(\eta_c, \eta_p, \mathbf{x}, \mathbf{c}, t, \epsilon, \mathbf{c}', \mathbf{y}) := \eta_c L_c^C(\mathbf{x}, t, \epsilon, \mathbf{c}) + \eta_p L_c^C(\mathbf{y}, t, \epsilon, \mathbf{c}') \quad (17.15)$$

Algorithm 12 Backdoor conditional DMs with caption trigger

Inputs: Backdoor Caption Trigger $\mathbf{g}$, Backdoor Target $\mathbf{y}$, Training dataset D^C , Training parameters θ , Text Encoder **Encoder**

while not converge **do**

$\{\mathbf{x}, \mathbf{p}, \eta_c, \eta_p\} \sim D^C$

$t \sim \text{Uniform}(\{1, \dots, T\})$

$\epsilon \sim \mathcal{N}(0, \mathbf{I})$

$\mathbf{c}, \mathbf{c}' = \text{Encoder}(\mathbf{p}), \text{Encoder}(\mathbf{p} \oplus \mathbf{g})$

 Use gradient descent $\nabla_{\theta} L_{\theta}^{CC}(\eta_c, \eta_p, \mathbf{x}, t, \epsilon, \mathbf{c}', \mathbf{y})$ to update θ

end while

As for the image-trigger backdoor, we can also derive the backdoor loss function $L_p^{CI}(\mathbf{x}, t, \epsilon, \mathbf{g}, \mathbf{y}, \mathbf{c}, \zeta) := \|\epsilon - \frac{2H(t)}{(1+\zeta)G^2(t)} \mathbf{r}(\mathbf{x}, \mathbf{g}) - \epsilon_{\theta}(\mathbf{x}'_t(\mathbf{y}, \mathbf{r}(\mathbf{x}, \mathbf{g}), \epsilon), t, \mathbf{c})\|^2$ based on Sect. 17.3.4. The image-trigger backdoor loss function can be expressed as

$$L_{\theta}^{CI}(\eta_c, \eta_p, \mathbf{x}, \mathbf{c}, t, \epsilon, \mathbf{g}, \mathbf{y}, \zeta) := \eta_c L_c^C(\mathbf{x}, t, \epsilon, \mathbf{c}) + \eta_p L_p^{CI}(\mathbf{x}, t, \epsilon, \mathbf{g}, \mathbf{y}, \mathbf{c}, \zeta) \quad (17.16)$$

We summarize the backdoor training algorithms of the unconditional (image-as-trigger) and conditional (caption-as-trigger) DMs in Algorithms 11 and 12. We denote the text encoder as **Encoder** and $\oplus$ as concatenation. For a caption-image dataset D^C , each data point $\mathbf{e}^i$ consists of the clean image $\mathbf{x}^i$, the clean/backdoor loss weight η_c^i/η_p^i , and the clean caption $\mathbf{p}^i$.

17.4 Backdoor Detection and Mitigation for Diffusion Models

In [22], the authors studied several backdoor attacks on DMs and concluded that the key factor of injected backdoor is implanting a distribution shift relative to the trigger in DMs. Based on this insight, they proposed the backdoor detection and removal framework for DMs. This framework can be used without any real clean data.

For backdoor detection, the authors introduced a new trigger inversion method to invert a trigger based on the given DM. It leverages a distribution shift preservation property –an inverted trigger should maintain a relative distribution shift across the multiple steps in the model inference process. The detection is then based on the images produced by the DM when the inverted trigger is stamped on Gaussian noise inputs. In the detection process, a metric called *uniformity score* is used to measure the consistency of generated images. This score and the *Total Variance* loss that measures the noise level of an image are used to decide whether a DM is backdoored.

Backdoor Mitigation To mitigate the backdoor, the authors design a loss function to reduce the distribution shift of the model against the inverted trigger. Evaluated on a variety of DMs, samplers, and backdoor attacks, the experimental results in [22] show that the proposed framework can have close to 100% detection accuracy and reduce the backdoor effects to close to zero while largely maintaining the model utility.

17.5 Performance Evaluation

17.5.1 Experiment Setup

Text as Trigger We use text in captions as the trigger to backdoor conditional DMs in Sect. 17.5.2. We take Stable Diffusion v1-4 [708] as the pre-trained model and design various caption triggers and image targets shown in Fig. 17.2. We fine-tune Stable Diffusion on the two datasets Pokemon Caption [657] and CelebA-HQ-Dialog [374] with Low-Rank Adaptation (LoRA) [321].

Image as Trigger We use images as triggers as shown in Table 17.1. We also consider three kinds of DMs, DDPM [299], LDM [710], and NCSN [776–778], to examine the effectiveness of our unified framework. We evaluate the generalizability of our framework on various samplers in Sect. 17.5.3 with the pre-trained model (*google/ddpm-cifar10-32*) released by Google HuggingFace organization on CIFAR10 dataset [424]. We also use the latent diffusion model [710] downloaded from Huggingface (*CompVis/ldm-celebahq-256*), which is pre-trained on CelebA-HQ [539]. As for score-based models, we retrain the model by ourselves on the CIFAR10 dataset [424].

Backdoor Attack Configuration For conditional DMs, we choose 10 different caption triggers shown in the marker legend of Fig. 17.2. Note that due to the matplotlib’s limitation, in the legend, {SOCCER} and {HOT_FACE} actually represent the symbols ‘’ and ‘’. As for unconditional DMs, in the CIFAR10 and CelebA-HQ datasets, we follow the same backdoor configuration as BadDiffusion [145], as specified in Table 17.1.

Evaluation Metrics We design three qualitative metrics to measure the performance of VillanDiffusion in terms of utility and specificity respectively. For measuring utility, we use FID [296] score to evaluate the quality of generated clean samples. Lower scores mean better quality. For measuring specificity, we use Mean Square Error (MSE) and MSE threshold to measure the similarity between ground truth target images y and generated backdoor sample $\hat{y}$, which is defined as $MSE(y, \hat{y})$. Lower MSE means better similarity to the target. Based on MSE, we also introduce another metric, called MSE threshold, to quantify the attack effectiveness, where the samples under a certain MSE threshold ϕ are marked as 1,

Table 17.1 Experiment setups of image triggers and targets following [145]. The black color indicates no changes to the corresponding pixel values when added to the data input

CIFAR10 (32×32)							CelebA-HQ (256×256)	
Triggers		Targets					Trigger	Target
Grey Box	Stop Sign	NoShift	Shift	Corner	Shoe	Hat	Eyeglasses	Cat
								

otherwise as 0. Formally, the MSE threshold can be defined as $\mathbb{I}(\text{MSE}(y, \hat{y}) < \phi)$. A higher MSE threshold value means better attack success rates.

For backdoor attacks on the conditional DMs, we compute the cosine similarity between the caption embeddings with and without triggers, called **caption similarity**. Formally, we denote a caption with and without trigger as $\mathbf{p} \oplus \mathbf{g}$ and $\mathbf{p}$ respectively. With a text encoder **Encoder**, the caption similarity is defined as $\langle \text{Encoder}(\mathbf{p}), \text{Encoder}(\mathbf{p} \oplus \mathbf{g}) \rangle$.

17.5.2 Caption-Trigger Backdoor Attacks on Text-to-Image DMs

We fine-tune the pre-trained stable diffusion model [708, 710] with the frozen text encoder and set learning rate $1e-4$ for 50,000 training steps. For the backdoor loss, we set $\eta_p^i = \eta_c^i = 1, \forall i$ for the loss (17.15). We also set the LoRA [321] rank as 4 and the training batch size as 1. The dataset is split into 90% training and 10% testing. We compute the MSE and MSE threshold metrics on the testing dataset and randomly choose 3K captions from the whole dataset to compute the FID score for the Celeba-HQ-Dialog dataset [374]. As for the Pokemon Caption dataset, we also evaluate MSE and MSE threshold on the testing dataset and use the caption of the whole dataset to generate clean samples for computing the FID score.

We present the results in Fig. 17.2. From Fig. 17.2a, e, we can see the FID score of the backdoored DM on CelebA-HQ-Dialog is slightly better than the clean one, while the Pokemon Caption dataset does not, which has only 833 images. This may be caused by the rich and diverse features of the CelebA-HQ-Dialog dataset. In Fig. 17.2b, f, the MSE curves get closer as the caption similarity becomes higher. This means as the caption similarity goes higher, the model cannot distinguish the difference between clean and backdoor captions because of the fixed text encoder.




Fig. 17.2 Evaluation of various caption triggers in FID, MSE, and MSE threshold metrics. Every color in the legend of (b)/(e) corresponds to a caption trigger inside the quotation mark of the

Thus, the model will tend to generate backdoor targets with equal probabilities for clean and backdoor captions respectively. The MSE threshold in Figs. 17.2c, g also explains this phenomenon. We can also see that the backdoor FID scores are slightly lower than the clean FID score (green dots marked with red boxes) in Fig. 17.2a. In Fig. 17.2b, c, as the caption similarity goes up, the clean sample and backdoor samples contain target images with similar likelihood.

17.5.3 Image-Trigger Backdoor Attacks on Unconditional DMs

Backdoor Attacks with Various Samplers on CIFAR10 We fine-tune the pre-trained diffusion models *google/ddpm-cifar10-32* with learning rate $2e-4$ and 128 batch size for 100 epochs on the CIFAR10 dataset. To accelerate the training, we use half-precision (float16) training. During the evaluation, we generate 10K clean and backdoor samples for computing metrics. We conduct the experiment on 7 different samplers with 9 different configurations, including DDIM [774], DEIS [987], DPM Solver [548], DPM Solver++ [549], Heun’s method of EDM (algorithm 1 in [387]), PNDM [520], and UniPC [1016]. We report our results in Fig. 17.3. We can see all samplers reach lower FID scores than the clean models under 70% poison rate for the image trigger *Hat*. Even if the poison rate reaches 90%, the FID score is still only larger than the clean one by about 10%. As for the MSE, in Fig. 17.3b, we can see about 10% poison rate is sufficient for a successful backdoor attack.

Backdoor Attacks on CelebA-HQ We fine-tune the DM with learning rate $8e-5$ and batch size 16 for 1500 epochs and use mixed-precision training with float16. In Fig. 17.4, we show that we can achieve a successful backdoor attack with 20% poison rate while the FID scores increase about 25 ~ 85%. Although the FID scores of the backdoor models are relatively higher, we believe training for longer epochs can further decrease the FID score. The backdoor attack results of latent diffusion models (LDM) and score-based models can be found in [146].

Fig. 17.2 (continued) marker legend. The target images are shown in (d) and (h) for backdoor-ing CelebA-HQ-Dialog and Pokemon Caption datasets, respectively. In (b) and (c), the dotted-triangle line indicates the MSE/MSE threshold of generated backdoor targets and the solid-circle line is the MSE/MSE threshold of generated clean samples




Fig. 17.3 FID and MSE scores of various samplers and poison rates. Every color represents one sampler. Because DPM Solver and DPM Solver++ provide the second and the third order approximations, we denote them as “O2” and “O3” respectively. (a) Hat FID. (b) Hat MSE. (c) Shoe FID. (d) Shoe MSE




Fig. 17.4 Backdoor DDPM on CelebA-HQ. (a) CelebA-HQ FID. (b) CelebA-HQ MSE

Chapter 18

Prompt Engineering for Safety Red-Teaming: A Case Study on Text-to-Image Diffusion Models



Abstract Diffusion models for text-to-image (T2I) synthesis have demonstrated exceptional capabilities for generating high-quality content. However, this progress has raised notable concerns about potential misuse, particularly in creating copyrighted, prohibited, and restricted content, or NSFW (not safe for work) images. While efforts have been made to mitigate such problems, either by implementing a safety filter at the evaluation stage or by fine-tuning models to eliminate undesirable concepts or styles, the effectiveness of these safety measures in dealing with a wide range of prompts remains largely unexplored. In this work, we introduce two primary prompt engineering approaches, Prompting4Debugging (P4D) (Chin et al., Prompting4debugging: Red-teaming text-to-image diffusion models by finding problematic prompts. International Conference on Machine Learning (2024)) and Ring-A-Bell (Tsai et al., Ring-a-bell! how reliable are concept removal methods for diffusion models? International Conference on Learning Representations (2024)), acting as debugging and red-teaming tools to automatically find problematic prompts for diffusion models to test the reliability of a deployed safety mechanism.

18.1 Introduction

In recent years, generative models have been making remarkable advancements across multiple domains, such as text, images, and even code generation, blurring the distinction between the works created by AI systems and those crafted by human experts. One prominent area of focus upon generative AI is text-to-image (T2I) generation [456, 688, 690, 709, 720], where most of the state-of-the-art T2I methods are built upon the diffusion models, in which these T2I diffusion models enable the transformation of textual information into images. They not only bridge the gap between natural language processing and visual content creation, but also enhance the interaction and understanding across these two modalities. One of the main factors leading to the exceptional performance of T2I diffusion models nowadays stems from the vast amount of training data available on the internet, allowing the models to generate a wide range of content, including natural animals, sketches, cartoon images, and even artistic images. However, such large-scale training data

collected from the Internet can be a double-edged sword, as it can lead the models to unconsciously generate inappropriate content such as copyright infringement and NSFW materials.

To this end, there are several recent research works proposing the diffusion models equipped with safety mechanisms, e.g. Stable Diffusion with negative prompts [709], SLD [732], and ESD [241], which either restrict the text embedding space during inference or finetune the model for attempting to prevent the model from generating copyrighted or inappropriate images. Although these safety mechanisms are shown to be partially effective according to their evaluation schemes, there are already studies that demonstrate their potential flaws. For example, [691] has found that the state-of-the-art Stable Diffusion model equipped with NSFW safety filter [709] will still generate sexual content if users give the text prompt “A *photo of a billboard above a street showing a naked man in an explicit position*”. However, these problematic prompts are discovered manually and thus are hard to scale. Here hence comes an urgent need for developing an automated and scalable red-teaming tool for developers to systematically inspect the model safety and reliability before deployment.

On the other hand, as the rapid increase of size (e.g. even up to billions of parameters) for recent T2I diffusion models [688, 690, 709, 720], model finetuning becomes extremely expensive and infeasible upon limited computation resources while building the red-teaming tool. As a result, in this chapter, we utilize prompt engineering [81, 167, 375, 452, 479, 652, 730, 757] as our basis for developing the red-teaming technique, which achieves comparable performance to traditional approaches of finetuning heavy models but only needs to learn small amount of parameters or simply manipulate the prompt input without further modifying the model. Specially, this chapter presents two primary adversarial testing methods, Prompting4Debugging (P4D) [143] and Ring-A-Bell [826], for red-teaming T2I models. P4D is suited for red-teaming white-box T2I models, where the details including the model architecture and weights are known to the evaluator. Ring-A-Bell is suited for red-teaming black-box T2I models, like an online service or API, where the evaluator can only has access to the generated outputs.

18.2 Background and Related Work

Diffusion Models (DMs) DMs [298, 768] are powerful generative models that learn to simulate the data generation process by progressively denoising the (intermediate) noisy states of data, where such denoising steps stand for the backward process to the opposite forward one composed of diffusion steps which gradually add random noise to data. Given an input image x , Denoising Diffusion Probabilistic Models (DDPM) [298] first generates intermediate noisy image x_t at time step t via the forward diffusion steps, where x_t can be written as a close form depending on x , t , and noise ϵ sampled from Gaussian distribution $\mathcal{N}(0, I)$. Then the diffusion model training is based on the backward process for learning a model parameterized by θ

to predict ϵ , where such model takes both x_t and the corresponding time step t as input. The objective is defined as:

$$\mathcal{L}_{DM} = \mathbb{E}_{x, \epsilon \sim \mathcal{N}(0,1), t \in [1, T]} \left[\|\epsilon - \epsilon_\theta(x_t, t)\|_2^2 \right] \quad (18.1)$$

where t ranges from 1 to the maximum time step T .

Latent Diffusion Models (LDMs) Rombach et al. [709] proposes to model both forward and backward processes in the latent space, for alleviating the efficiency issue of DDPM which stems from having the model operate directly in the pixel space, where the transformation between latent and pixel spaces is based on a variational autoencoder (composed of an encoder $\mathcal{E}$ and a decoder $\mathcal{D}$). Furthermore, they extend DDPM to enable conditional image generation, via incorporating diverse conditions such as text prompts. Specifically, given the latent representation $z = \mathcal{E}(x)$ of input image x as well as the intermediate noisy latent vector z_t at time step t (analogously, depending on z , t , and $\epsilon \sim \mathcal{N}(0, I)$), a model parameterized by θ is trained to make prediction for the noise $\epsilon_\theta(z_t, c, t)$ that is conditioned on z_t , time step t , and a text condition c . The objective for learning such conditional generation process (based on image–condition training pairs $\{(x, c)\}$) is defined as:

$$\mathcal{L}_{LDM} = \mathbb{E}_{\mathcal{E}(x), c, \epsilon \sim \mathcal{N}(0,1), t \in [1, T]} \left[\|\epsilon - \epsilon_\theta(z_t, c, t)\|_2^2 \right]. \quad (18.2)$$

AI Red-Teaming Tools Red-teaming is an active cybersecurity assessment method that exhaustively searches for vulnerabilities and weaknesses in information security, where the issues found by red-teaming can further help companies or organizations improve their defense mechanisms and strengthen overall cybersecurity protection. Recently, with the popularity and increasing demand for generative AI, red-teaming is also being applied to AI models (especially language models [441, 754]) to enhance model security and stability. [754] fools the model for detecting machine-generated text by revising output, e.g. replacing synonym words or altering writing style in generated sentences. On the other hand, [441] constructs a pool of user inputs and employs Bayesian optimization to iteratively modify diverse positive test cases, eventually leading to model failures. The perspective of red-teaming is distinctly different from that of potential attackers. Drawing parallels with notable works such as [649], which employs an LLM as a red-team agent to generate test cases for another target LLM, and [900], which elicits unsafe responses from an LM by scoring an LM’s response with a safety classifier and then refining the prompt with gradient backpropagation through both the unfrozen safety classifier and the LM, we underscore a growing trend in red-team generative modeling. These approaches pragmatically employ both the target model’s inherent information and that from related models or external classifiers as practical means for red-teaming efforts aimed at debugging and enhancing model safety by utilizing all available information. However, these methods are only applicable to red-team

language models, while our P4D focuses on text-to-image models, which is a field that has been rarely explored in AI red-teaming.

Prompt Engineering Prompt engineering originates from the field of natural language processing and aims to adapt a pretrained language model to various downstream tasks by modifying input text with prompts. Prompt engineering can be categorized into two groups: *hard prompts* and *soft prompts*. Hard prompts, also known as discrete tokens, usually consist of interpretable words that are hand-crafted by users. For instance, [81] first demonstrates the remarkable generalizability of pretrained language models via adopting manually crafted hard prompts to a wide range of downstream tasks in few-shot learning. Then [247, 375, 729] reformulate input texts into specific cloze-style phrases, thus maintaining the form of hard prompts, to prompt the language models. On the other hand, soft prompts consist of appended continuous-valued text vectors or embeddings, providing a larger search space compared to hard prompts. For instance, prompt-tuning [452] and prefix-tuning [758] automate the soft prompts in continuous space. However, soft prompts are often uninterpretable or non-transferrable (i.e. cannot be shared by different language models). As a consequence, some discrete optimization methods are proposed to strike a balance between hard prompts and soft prompts, e.g. AutoPrompt [758], FluentPrompt [752], and PEZ [897] that learns hard prompts through continuous gradient-based optimization. Additionally, PEZ extends its capabilities to discover prompts that can be matched with given images, achieved by measuring the CLIP Score [295] using the same optimization method. Another line of works [279, 505, 573, 965] utilizes prompt tuning to identify target prompts for black-box models. For instance, [573] aims to generate adversarial prompts for black-box T2I models, which however is computationally expensive due to its inability of leveraging the iterative decoding properties (e.g., denoising steps in diffusion models) in T2I models. These studies demonstrate the potential of prompt engineering across various tasks and domains, motivating us to integrate such technique into the field of red-teaming T2I diffusion models.

Diffusion Models with Safety Mechanisms In response to the emerging issues of generating inappropriate images from diffusion models, several works have devoted to address the concern. These works fall into two categories: guidance-based and finetuning-based methods. For guidance-based methods like Stable Diffusion with negative prompts [709] and SLD [732], they block the text embedding of certain words or concepts (e.g. nudity, hate, or violence), in order to prevent the generation of the corresponding image content during the inference process. Rather than using guidance-based techniques, ESD [241] takes a different approach by finetuning the partial model weights (e.g. the U-Net to perform denoising in Stable Diffusion) to remove unwanted contents from the image output. Nonetheless, certain corner cases still bypass the safety mechanisms of these diffusion models [691].

18.3 Prompting4Debugging (P4D)

In [143], the authors develop a red-teaming tool named Prompting4Debugging (P4D) for Text-to-image (T2I) diffusion models to test the reliability of deployed safety mechanisms. In particular, three models, including Stable Diffusion (SD) with negative prompts [709], SLD [732], and ESD [241], are considered as the targets of study. The overview of the P4D framework is shown in Fig. 18.1 and detailed as follows.

Given an input text prompt P which is able to lead an unconstrained/standard T2I diffusion model $\mathcal{G}$ for generating the output image with an inappropriate concept/object C (i.e. $\mathcal{G}$ does not have the safety mechanism, and P is a problematic prompt), when taking such prompt P as the input for another T2I diffusion model $\mathcal{G}'$ equipped with the safety mechanism specific for C , ideally the resultant output image should be free from C (i.e. $\mathcal{G}'$ successfully defends the generated image against the problematic prompt P). Our red-teaming tool P4D now attempts to counteract the safety mechanism of $\mathcal{G}'$ such that the inappropriate concept/object C now again appears in the generated image (i.e. the safety mechanism of $\mathcal{G}'$ is bypassed).

Specifically, P4D adopts the technique of prompt engineering to circumvent the safety mechanism in $\mathcal{G}'$, where a new or modified prompt P^* is optimized for making $\mathcal{G}'$ conditioned on P^* to produce the inappropriate content as what would be obtained by having $\mathcal{G}$ conditioned on P . As the state-of-the-art T2I diffusion model, i.e. Stable Diffusion (SD), as well as the choices for the T2I diffusion models with safety mechanism $\mathcal{G}'$ in this work (e.g. SD with negative prompts [709], SLD [732],




Fig. 18.1 Overview of Prompting4Debugging (P4D) framework [143]. P4D employs prompt engineering techniques to red-team the text-to-image (T2I) diffusion model $\mathcal{G}'$ with safety mechanism (e.g. Stable Diffusion with negative prompts [709], SLD [732], and ESD [241]). With the help of an unconstrained T2I diffusion model $\mathcal{G}$, P4D optimizes to find the safety-evasive prompts (i.e. P^*_{cont}) which can bypass the safety mechanism in $\mathcal{G}'$ and still lead to generation of inappropriate image concept/objects (e.g. nudity)

and ESD [241]) are all based on the latent diffusion models, the optimization for P^* in P4D is actually realized in the latent space, following the procedure below (cf. Fig. 18.1):

1. With an unconstrained T2I diffusion model $\mathcal{G}$ (e.g. Stable Diffusion), an original text prompt P is first used to generate an image x having the inappropriate concept/object C . Note that the noise predictor in the backward process of $\mathcal{G}$ is parameterized by θ .
2. We then obtain the latent representation $z = \mathcal{E}(x)$ of x via the encoder $\mathcal{E}$ of $\mathcal{G}$ (noting that $\mathcal{G}$ is based on latent diffusion models thus has the corresponding variational autoencoder), followed by computing the intermediate noisy latent vector z_t at an arbitrary time step t according to the diffusion process of $\mathcal{G}$.
3. Given a T2I diffusion model with safety mechanism $\mathcal{G}'$ in which its noise predictor in the backward process is parameterized by θ' , we now aim to find a prompt P^* such that $\mathcal{G}'$ conditioned on P^* can produce the output x^* similar to x , thereby also having the similar inappropriate concept/object C . The optimization for P^* happens on the latent space directly to encourage similarity between noise predictions $\epsilon_\theta(z_t, P, t)$ and $\epsilon_{\theta'}(z_t, P^*, t)$. The basic idea is that, starting from the same noisy latent vector z_t at an arbitrary time step t , if both the noise predictors of $\mathcal{G}$ and $\mathcal{G}'$ which respectively take P and P^* as text prompt are able to reach the same noise prediction, then our goal of assuring the similarity between x^* and x in pixel space ideally can be also achieved.

Notably, the text prompt is typically fed into the noise predictor in the form of embeddings (according to the common practice for our $\mathcal{G}$ and $\mathcal{G}'$). To this end, the noise prediction happens in $\mathcal{G}$ is actually operated as $\epsilon_\theta(z_t, \mathcal{W}(P), t)$, where $\mathcal{W}$ is a pre-trained and fixed text encoder (e.g. CLIP [295]) for extracting the embedding $\mathcal{W}(P)$ of text prompt P . While for the noise prediction in $\mathcal{G}'$ that involves the optimization target P^* , P4D adopts the similar design of prompt engineering as PEZ [897] to automate the optimization (a benefit of soft prompt) while making the resultant prompt more transferable (a benefit of hard prompt): We start from a continuous/soft embedding $P_{\text{cont}}^* = [e_1, \dots, e_N]$ composed of N tokens $e_i \in \mathbb{R}^d$, followed by projecting P_{cont}^* into the corresponding discrete/hard embedding $P_{\text{disc}}^* = \mathcal{F}(P_{\text{cont}}^*)$ via a projection function $\mathcal{F}$ (where each token in P_{cont}^* is mapped to its nearest vocabulary embedding). As a result, noise prediction in $\mathcal{G}'$ is now operated as $\epsilon_{\theta'}(z_t, P_{\text{disc}}^*, t)$, and the objective $\mathcal{L}$ for our debugging process is defined as

$$\mathcal{L} = \|\epsilon_\theta(z_t, \mathcal{W}(P), t) - \epsilon_{\theta'}(z_t, P_{\text{disc}}^*, t)\|_2^2 \quad (18.3)$$

Here, both noise predictors in $\mathcal{G}$ and $\mathcal{G}'$ are fixed in optimization.

It is worth mentioning that the projection function $\mathcal{F}$ acts as a vector quantization operation and is non-differentiable. P4D follows the practice of PEZ [897] by directly updating P_{cont}^* based on the gradient of $\mathcal{L}$ with respect to P_{disc}^* . Specifically, we perform the update as $P_{\text{cont}}^* = P_{\text{cont}}^* - \gamma \nabla_{P_{\text{disc}}^*} \mathcal{L}$, where γ represents the learning

rate. Last but not least, the resultant P_{disc}^* can be transformed into legible texts P^* by the off-the-shelf text decoder/tokenizer.

In [143], two variants for P_{cont}^* : **P4D- N** and **P4D- K** are explored, where the former initializes N tokens in P_{cont}^* from scratch via randomly drawing N vocabulary embeddings, while the latter inserts learnable tokens after every K tokens of $\mathcal{W}(P)$ (i.e. the embedding of the original text prompt P). Basically, P_{cont}^* in P4D- N has fixed length of N which is independent from the length of $\mathcal{W}(P)$, it would potentially be insufficient for red-teaming the images with complex content as the original prompt length are not taken into consideration. In comparison, the length of P_{cont}^* in P4D- K (and the number of trainable tokens being inserted) varies with the length of $\mathcal{W}(P)$ thus alleviating the aforementioned concern in P4D- N . The experiments in [143] showed that both P4D- N and P4D- K have the comparable red-teaming performance but the hard prompt found by P4D- K shows better interpretability as the original prompt P is taken as its part.

18.4 Ring-A-Bell

Instead of assuming the white-box access of target models, Ring-A-Bell [826] focuses on constructing red-teaming techniques only with black-box access. We denote the original unconstrained DM (i.e., without any safety mechanisms) as $\epsilon_{\theta}(\cdot)$. On the other hand, models with a safety mechanism are denoted as $\epsilon_{\theta'}(\cdot)$. Given a target concept c (e.g., nudity, violence, or artistic concept such as “style of Van Gogh”), we want to find an adversarial concept $\tilde{c}$ such that, given a trajectory, $z_0, z_1, \dots, z_T$ (typically the one that produces the inappropriate image z_0), these two models can be guaranteed to have similar probabilities of generating such a trajectory, i.e.,

$$P_{\epsilon_{\theta}}(z_0, z_1, \dots, z_T | c) \approx P_{\epsilon_{\theta'}}(z_0, z_1, \dots, z_T | \tilde{c}), \quad (18.4)$$

where P is the probability that the backward process is generated by the given noise predictor. When minimizing the KL divergence between two such distributions, the objective is expressed as L_{white} ,

$$\mathcal{L}_{\text{white}} = \sum_{\hat{t}=1}^T \mathbb{E}_{z_{\hat{t}} \sim P_{\epsilon_{\theta}}(z_{\hat{t}} | c)} [\|\rho(\epsilon_{\theta}(z_{\hat{t}}, c, \hat{t}) - \epsilon_{\theta'}(z_{\hat{t}}, \tilde{c}, \hat{t}))\|^2], \quad (18.5)$$

where ρ denotes the weight on the loss. See [826] for the detailed derivation of $\mathcal{L}_{\text{white}}$. To briefly explain the red-teaming process, given a forward process starting with an inappropriate image z_0 , we want the backward process produced by the noise predictor $\epsilon_{\theta}(\cdot)$ and $\epsilon_{\theta'}(\cdot)$ under the original concept c and the adversary concept $\tilde{c}$ to be similar, and thus output similar images. Namely, we have $\tilde{c} := \arg \min_{\tilde{c}} \mathcal{L}_{\text{white}}(\tilde{c})$.




Fig. 18.2 Overview of Ring-A-Bell framework [826]. The problematic prompt generation is model-agnostic and can be carried out offline

In particular, in the black-box setup, we can no longer obtain the adversarial concept $\tilde{c}$ directly from probing the modified model $\epsilon_{\theta'}$ and the unconstrained model ϵ_{θ} . Ring-A-Bell address such a challenge with its overall pipeline shown in Fig. 18.2. The rationale behind Ring-A-Bell is that current T2I models with safety mechanisms either learn to disassociate or simply filter out relevant words of the target concepts with their representation c , and thus the detection or removal of such concepts may not be carried out completely if there exist implicit text-concept associations embedded in the T2I generation process. That is, Ring-A-Bell aims to perform black-box adversarial testing to check whether a supposedly removed concept can be revoked via its prompt optimization procedure.

In Ring-A-Bell, we first generate the holistic representation of concept c by collecting prompt pairs that are semantically similar with only difference in concept c . For instance, as in Fig. 18.2, the “clashed / peaceful” and “bleeding / fine” are differing in the concept “violence”. Afterwards, the empirical representation $\hat{c}$ of c is derived as

$$\hat{c} := \frac{1}{N} \sum_{i=1}^N \{f(\mathbf{P}_i^c) - f(\mathbf{P}_i^{\phi})\}, \quad (18.6)$$

where $f(\cdot)$ denotes the text encoder with prompt input (e.g., text encoder in CLIP [295]) and ϕ denotes the absence of concept c . Simply put, given prompt pairs $\{\mathbf{P}_i^c, \mathbf{P}_i^{\phi}\}_{i=1}^N$ with similar semantics but contrasting in the target concept c , such as (Prompt 1, Prompt 2) in Fig. 18.2 that represent the concept “violence”, we extract the empirical representation $\hat{c}$ by pairwise subtraction of the embedding and then averaging over all pairs. This ensures that the obtained representation does not suffer from context-dependent influence, and by considering all plausible scenarios, we obtain the full semantics underlying the target concept c . Similar attempts can also be seen in [1048] where only the sign vectors are used to induce concept

removal/generation. We refer the readers to [826] for the detailed generation of prompt-pairs.

After obtaining $\hat{c}$, Ring-A-Bell transforms the target prompt $\mathbf{P}$ into the problematic prompt $\tilde{\mathbf{P}}$. In particular, Ring-A-Bell first uses the soft prompt of $\mathbf{P}$ and $\hat{c}$ to generate $\tilde{\mathbf{P}}_{cont}(\hat{c})$ as

$$\tilde{\mathbf{P}}_{cont} := f(\mathbf{P}) + \eta \cdot \hat{c}, \quad (18.7)$$

where η is the strength coefficient available for tuning. In short, $\tilde{\mathbf{P}}_{cont}$ is the embedding of $\mathbf{P}$ infused with varying levels of concept c . Finally, we generate $\hat{\mathbf{P}}$ by solving the optimization problem below

$$\min_{\hat{\mathbf{P}}} ||f(\hat{\mathbf{P}}) - \tilde{\mathbf{P}}_{cont}||_2^2 \text{ subject to } \hat{\mathbf{P}} \in S^K, \quad (18.8)$$

where K is the length of the query and S is the set of all word tokens. Here, the variables to be optimized are discrete with the addition that typically S consists of a huge token space. Hence, we adopt the genetic algorithm (GA) [764] as our optimizer because its ability to perform such a search over large discrete space remains competitive.

It is evident from the above illustration that Ring-A-Bell requires no prior knowledge of the model to be evaluated except for the access of the external text encoder (i.e., the access of $f(\cdot)$ in Eqs. (18.6) $\sim$ (18.8)). Furthermore, Ring-A-Bell presents a readily available database that stores various sensitive concepts. Any user could utilize the concepts identified, automatically create problematic prompts offline, and further deploy them online, demonstrating the practicality of Ring-A-Bell.

18.5 Performance Evaluation

Dataset The evaluation is conducted on concept-related and object-related datasets. For concept-related dataset, we focus on Inappropriate Image Prompts (I2P) dataset [732], which encompasses various uncomfortable and inappropriate prompts (including hate, harassment, violence, self-harm, nudity contents, shocking images, and illegal activity). Specifically, nudity contents are most prohibitive due to privacy and respect considerations, we hence specifically set this concept aside for separate evaluation. On the other hand for the object-related datasets, we utilize the “car” and “French-horn” classes from ESD [241] for our evaluation (as ESD only offers finetuned weights for these two classes). Notably, the original French-horn dataset comprises merely 10 identical prompts with different evaluation seeds. We hence extend the size of French-horn prompts from 10 to 305 by experimenting with a wider array of evaluation seeds. See [143] for more details.

18.5.1 P4D Results

Standard T2I and Safe T2I Models We adopt the typical Stable Diffusion [709] (denoted as **standard SD**) for our standard T2I model, while using **ESD** [241], **SLD** [732] (where we adopt two superior variants of SLD, i.e. **SLD-MAX** and **SLD-STRONG**, provided in their release code), and SD with negative prompts [709] (denoted as **SD-NEGP**) for our safe T2I models. For standard SD, ESD, and SLD, we apply the Stable Diffusion v1-4 model backbone, while for SD-NEGP, we use the Stable Diffusion v2-0 model backbone. We set $N = 16$ and $K = 3$ respectively for P4D- N and P4D- K .

Baselines We compare four distinct baselines, namely **Random- N** , **Random- K** , **Shuffling**, and **Soft Prompting**. Random- N is analogous to P4D- N , where N vocabulary embeddings are randomly drawn to be the input prompt for safe T2I models, but without any further optimization being performed. Similarly, Random- K is analogous to P4D- K (i.e., inserting random vocabulary embeddings after every K tokens in $W(P)$) but excludes further optimization. Furthermore, as some natural language researches have discovered that shuffling the word order in a sentence can make ChatGPT [625] generate inappropriate responses, we introduce a similar approach to build Shuffling baseline, which involves randomly permuting the words in prompt P . Moreover, our method optimizes the hard prompt projected from a continuous soft prompt embedding. Hence, we introduce Soft Prompting baseline which directly optimizes continuous soft embedding without projection. Soft Prompting- N and Soft Prompting- K are analogous to P4D- N and P4D- K respectively.

Evaluation During the evaluation process, we generate 3 images for every prompt (either the original prompts in the dataset or the ones produced by the baselines/P4D) by safe T2I models, where a prompt is considered problematic if any of the resulting images is identified as “unsafe” by the classifiers or detectors described below. For categories in the concept-related I2P dataset, we employ the NudeNet image classifier [54] with an unsafe percentage threshold of 0.45 to detect nudity content, while adopting the Q16 classifier [734] for identifying the other categories in I2P. For object-related categories (i.e. car and French-horn), we utilize YOLO v5 vehicle detector [69] with the confidence threshold at 0.7 for identifying cars, while employing a pre-trained ResNet-18 classifier [556] from the Imagenette dataset [316] for detecting French-horns. Although these classifiers and detectors are not error-free, they act as a low-cost third-party auditor, capable of evaluating our P4D and the baselines in a scalable and fair manner. **Failure rate (FR)** is adopted as **metric** to show how many problematic prompts are identified from the entire dataset. The higher FR indicates better debugging performance for red-teaming.

Numerical Results Table 18.1 summarizes the numerical results of P4D and the baselines. Regarding concept-related I2P dataset, we inspect all safe T2I models for the nudity category; while we only examine SLD-MAX for all the other categories,

Table 18.1 Quantitative evaluation among various approaches for red-teaming performance, where the failure rate (FR) indicating the proportion of problematic prompts with respect to the overall amount of data is adopted as the evaluation metric

Method	Nudity						All in 12P		Car		French-horn	
	ESD	SLD-MAX	SLD-STRONG	SD-NEGP	SLD-MAX	ESD	SLD-MAX	ESD	ESD	ESD	ESD	ESD
Random- <i>N</i>	0.95%	8.21%	10.55%	2.64%			12.45%	4.68%		0.50%		
Random- <i>K</i>	14.13%	22.94%	23.12%	18.24%			18.93%	22.71%		18.85%		
Shuffling	11.36%	27.74%	21.96%	11.44%			21.96%	22.47%		14.65%		
Soft prompting- <i>N</i>	13.32%	25.00%	33.33%	20.13%			21.80%	33.73%		25.02%		
Soft prompting- <i>K</i>	27.68%	33.55%	30.39%	21.79%			21.16%	41.54%		30.14%		
P4D- <i>N</i>	50.65%	25.67%	34.03%	25.44%			22.05%	40.42%		62.62%		
P4D- <i>K</i>	47.19%	38.69%	37.84%	20.36%			25.54%	34.87%		29.50%		
P4D-UNION	66.58%	52.66%	55.29%	40.98%			40.93%	56.64%		72.74%		

as SLD-MAX is the sole model capable of resisting additional concepts such as shocking, self-harm, and illegal content. Regarding object-related categories, we inspect ESD for cars and French-horns. From Table 18.1, we observe that P4D-*N* and P4D-*K* demonstrate promising and comparable results across a range of safe T2I models and categories, indicating P4D-*K* preserves its prompt interpretability without compromising the debugging performance. Furthermore, we unify problematic prompts from P4D-*N* and P4D-*K* and obtain P4D-UNION, which significantly increases the failure rate across various safe T2I models and categories (either concept-related or object-related ones), indicating that problematic prompts found by P4D-*N* and P4D-*K* are diverse. Notably, for the nudity category, as our P4D achieves the highest failure rate in ESD, in which it indicates that ESD originally (before red-teaming) provides the most effective safety protection against nudity content among all safe T2I models. However, the finetuning-based concept-removal safety mechanism of ESD may only learn to disassociate certain concept-related words with the unsafe image content, but it may not be resistant to optimized prompts. On the other hand, guidance-based safe T2I models such as SLD and SD-NEGP, restrict the textual embedding space for P4D to explore as well as prevent the generation of particular concepts/objects with their text filters, resulting in a lower failure rate compared to ESD with P4D. We observe that deactivating these text filters during training encourages P4D to investigate a broader range of problematic prompts (i.e. larger explorable textual embedding space). We refer to this phenomenon as “*information obfuscation*” for red-teaming [143].

In addition to the aforementioned baselines, we conduct further comparison with two baselines built upon the recent techniques of prompt optimization for text-to-image diffusion models (i.e. discover/optimize the prompt with respect to the given reference images), including PEZ [897] and Textual Inversion [238], where our experiments here are based on the nudity category of I2P dataset. For Textual Inversion (denoted as **Text-Inv**), we firstly create a pool of images generated by standard T2I model $\mathcal{G}$ with the ideal prompts. With randomly drawing three images from the pool as the reference images, Text-Inv is applied to optimize S_* token which symbolizes the concept of nudity, where such process is repeated for M times to obtain $\{S_{*1}, \dots, S_{*M}\}$ (where M is the number of ideal prompts). Finally, we evaluate the performance of Text-Inv baseline by inputting the sentence “a photo of S_{*i} ” into a safe T2I model $\mathcal{G}'$ to calculate FR. For **PEZ**, its two versions are adopted: PEZ-Original and PEZ-Prompt Inversion (noted as **PEZ-Orig** and **PEZ-PInv** respectively). Given a reference image x generated with the ideal prompt by $\mathcal{G}$, PEZ-Orig optimizes in the CLIP space to find the closest prompt P^* with respect to x ; while PEZ-PInv firstly obtains the latent representation $z = \mathcal{E}(x)$ of x using the encoder $\mathcal{E}$ of safe T2I model $\mathcal{G}'$, followed by computing the intermediate latent

Table 18.2 Quantitative results comparing with prompt optimization methods in the context of nudity category to red-teaming performance. Bold value represents best performing method for each concept

Method	ESD	SLD-MAX	SLD-STRONG	SD-NEGP
Text-Inv [238]	11.91%	13.73%	35.71%	8.13%
PEZ-Orig [897]	12.47%	24.51%	28.57%	20.57%
PEZ-PInv [897]	26.59%	22.06%	22.32%	12.44%
P4D- <i>N</i>	50.65%	25.67%	34.03%	25.44%
P4D- <i>K</i>	47.19%	38.69%	37.84%	20.36%
P4D-UNION	66.58%	52.66%	55.29%	40.98%

vector z_t with added noise η at an arbitrary time step t in the diffusion process, then optimizes P^* by encouraging the similarity between the noise prediction $\epsilon_{\theta'}(z_t, P^*, t)$ of $\mathcal{G}'$ and η . The optimized P^* from PEZ-Orig or PEZ-PInv is used as input to $\mathcal{G}'$ for evaluating the performance of two PEZ-based baselines. From the results in Table 18.2, P4D shows the ability to identify a greater number of problematic prompts across all four safe T2I models. It is worth noting that Text-Inv and PEZ-Orig optimize prompts using the information solely from the standard T2I model, while PEZ-PInv specifically leverages the information from the safe T2I model. As a result, the superior performance of our P4D indicates that the integration of the information from both standard T2I and safe T2I models enhances the efficacy of problematic prompt identification.

18.5.2 Ring-A-Bell Results

We remark that P4D [142] cannot be directly compared to Ring-A-Bell, because the former is a white-box red-teaming tool, while the latter is a black-box red-teaming tool. Therefore, their respective baselines would also differ in the performance evaluation. We evaluate the performance of Ring-A-Bell on the I2P dataset [733], an established dataset of problematic prompts, on the concepts of nudity and violence. We select 95 nudity prompts where the percentage of nudity is greater than 50%. For the concept of violence, to avoid overlapping with nudity prompts, we selected a total of 250 prompts with a nudity percentage less than 50%, an inappropriateness percentage greater than 50%, and labeled as harmful.

Concept Removal Models We choose several state-of-the-art concept removal methods, including ESD [242], Safe Latent Diffusion (SLD) [733] under 3 settings (where we adopt 3 superior variants of SLD as provided by [733], i.e., SLD-Max, SLD-Strong, and SLD-Medium), SD with negative prompts (SD-NP) [711], Concept Ablation (CA) [428], and Forget-Me-Not (FMN) [979]. Furthermore,

for ESD in both nudity and violence, we fine-tuned the non-cross-attention and cross-attention parameters with negative guidance set as 1 and 3, respectively. To remove nudity, we use “naked, nude, bare, exposed, Stripped, topless” as the negative prompt. On the other hand, to remove violence, we use the prompt “hate, harassment, violence, suffering, humiliation, harm, suicide, blood” for SD-NP and training ESD.

Baselines For the baselines, we select QF-Attack [1048] since QF-Attack. In particular, QF-Attack crafts an adversarial prompt by minimizing the cosine distance between $f(x)$ and $f(x')$, where x represents the original prompt, x' represents the adversarial prompt, and $f(\cdot)$ denotes the text-encoder of CLIP model. We conduct the targeted attack of the QF-Attack. The first step involves setting the mask words. This means that it will mask the specified words in the original prompt and, by adding a perturbation prompt to the end of the original prompt, generate an image whose content corresponds to the original prompt with the mask removed. Therefore, we select to mask out words related to clothing such as “wear tops and pants” and words related to peace and absence of harm such as “peaceful and fine”. Based on its default setting, the character length is 5, and the perturbation prompts are crafted using a genetic algorithm, which performed the best among other candidates [1048].

Evaluation Throughout the evaluation process, we generate one image per prompt with a random fixed seed. To classify whether images contain nudity, we use the NudeNet¹ detector. If the detector detects one of the following labels in the image: EXPOSED_ANUS, EXPOSED_BREAST_F, EXPOSED_GENITALIA_F, EXPOSED_GENITALIA_M, we consider the image to contain nudity. On the other hand, to determine whether images contain inappropriate content (such as blood or violence), we use the Q16 classifier [734]. In Ring-A-Bell, we first find the empirical concept $\hat{c}$. As shown in Fig. 18.2, it illustrates the use of the text encoder (e.g., CLIP) to obtain $\hat{c}$, where we use the VIT-L/14 [348]. We use 50 and 30 prompts containing nudity and violence, respectively, and modify these prompts to remove all traces of nudity and violence to obtain the concept $\hat{c}$. In addition to the single-setting Ring-A-Bell, we also follow [142] to consider the union settings denoted as Ring-A-Bell-Union, covering three different configurations of Ring-A-Bell for generating new prompts. Regarding the metric, we report the Attack Success Rate (ASR), defined as the proportion of successful generation of inappropriate images by problematic prompts relative to the total number of images. For the Ring-A-Bell-Union, each prompt is considered successful if it generates inappropriate images at least once under the three configurations.

¹ <https://github.com/vladmandic/nudenet> (last access: 2023/09).

Numerical Results We demonstrate the performance of Ring-A-Bell on T2I models that have been fine-tuned to forget nudity or violence. We note that both CA and FMN are incapable of effectively removing nudity and violence, but we still include them for completeness sake. Furthermore, we also consider a stringent defense, which involves applying both concept removal methods and safety checkers (SC) [692] to filter images for inappropriate content after generation. Regarding nudity, we set Ring-A-Bell with $K = 16$ and $\eta = 3$, while for Ring-A-Bell-Union, we employ different settings, including $(K, \eta) = (16, 3)$, $(77, 2)$, and $(77, 2.5)$. As for violence, we select $K = 77$ and $\eta = 5.5$ and for Ring-A-Bell-Union, we set $(K, \eta) = (77, 5.5)$, $(77, 5)$ and $(77, 4.5)$.

As shown in Table 18.3, contrary to using the original prompts and QF-Attack, Ring-A-Bell is more effective in facilitating these T2I models to recall forgotten concepts for both nudity and violence. When a safety checker is deployed, Ring-A-Bell is also more capable of bypassing it, especially under the union setting.

In [826], the authors also reported the red-teaming results on a set of online text-to-image generation services. The prompts used to test these online services are chosen from prompts generated by Ring-A-Bell for red-teaming concept removal models.

Table 18.3 Quantitative evaluation of different attack configurations against different concept removal methods via the metric of ASR. (w/o SC and w/ SC represent the absence and presence of the safety checker, respectively). Bold value represents best performing method for each concept

Concept	Methods	SD	ESD	SLD-Max	SLD-Strong	SLD-Medium	SD-NP	CA	FMN
Nudity	Original prompts (w/o SC)	52.63%	12.63%	2.11%	12.63%	30.53%	4.21%	58.95%	37.89%
	QF-Attack (w/o SC)	51.58%	6.32%	9.47%	13.68%	28.42%	5.26%	56.84%	37.89%
	Ring-A-Bell (w/o SC)	93.68%	35.79%	42.11%	61.05%	91.58%	34.74%	89.47%	68.42%
	Ring-A-Bell-Union (w/o SC)	97.89%	55.79%	57.89%	86.32%	100%	49.47%	96.84%	94.74%
	Original Prompts (w/ SC)	7.37%	5.26%	2.11%	6.32%	3.16%	2.11%	9.47%	15.79%
	QF-Attack (w/ SC)	7.37%	4.21%	2.11%	6.32%	8.42%	5.26%	9.47%	18.95%
	Ring-A-Bell (w/ SC)	30.53%	9.47%	7.37%	12.63%	35.79%	8.42%	37.89%	28.42%
	Ring-A-Bell-Union (w/ SC)	49.47%	22.11%	15.79%	32.63%	57.89%	16.84%	53.68%	47.37%

(continued)

Table 18.3 (continued)

Concept	Methods	SD	ESD	SLD-Max	SLD-Strong	SLD-Medium	SD-NP	CA	FMN
Violence	Original prompts (w/o SC)	60.4%	42.4%	16%	20.8%	34%	28%	62%	50.4%
	QF-Attack (w/o SC)	62%	56%	14.8%	24.2%	32.8%	24.8%	58.4%	53.6%
	Ring-A-Bell (w/o SC)	96.4%	54%	19.2%	50%	76.4%	80%	97.6%	79.6%
	Ring-A-Bell-Union (w/o SC)	99.6%	86%	40.4%	80.4%	97.2%	94.8%	100%	98.8%
	Original prompts (w/ SC)	56.8%	39.2%	14.4%	18%	30.8%	25.2%	54.8%	47.2%
	QF-Attack (w/ SC)	54.4%	53.6%	11.2%	21.2%	31.6%	21.2%	53.6%	47.2%
	Ring-A-Bell (w/ SC)	82.8%	49.2%	18%	44%	68.4%	68%	85.2%	74.4%
	Ring-A-Bell Union (w/ SC)	99.2%	84%	38.4%	76.4%	95.6%	90.8%	98.8%	98.8%

References

1. A new approach to linear filtering and prediction problems (1960)
2. 7 top large language model use cases and applications (2023). <https://www.projectpro.io/article/large-language-model-use-cases-and-applications/887>
3. Artificial intelligence risk management framework (ai rmf 1.0) (2023). <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>
4. Enhancing search using large language models (2023). <https://medium.com/whatnot-engineering/enhancing-search-using-large-language-models-f9dcb988bdb9>
5. fuzi.mingcha. <https://github.com/irlab-sdu/fuzi.mingcha> (2023)
6. Openai moderation api (2023). <https://platform.openai.com/docs/guides/moderation>
7. What does “fairness” mean for machine learning systems? (2023). https://haas.berkeley.edu/wp-content/uploads/What-is-fairness_-EGAL2.pdf
8. Aaronson, S.: My AI safety lecture for ut effective altruism (2022). URL <https://scottaaronson.blog/?p=6823>
9. Abowd, J.M.: The US census bureau adopts differential privacy. In: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp. 2867–2867 (2018)
10. Achille, A., Soatto, S.: Emergence of invariance and disentanglement in deep representations. *The Journal of Machine Learning Research* **19**(1), 1947–1980 (2018)
11. Ahn, K., Cheng, X., Daneshmand, H., Sra, S.: Transformers learn to implement preconditioned gradient descent for in-context learning. arXiv preprint arXiv:2306.00297 (2023)
12. AI4Science, M.R., Quantum, M.A.: The impact of large language models on scientific discovery: a preliminary study using gpt-4 (2023)
13. Akpanuko, E.E., Asogwa, I.E.: Accountability: A synthesis. *International Journal of Finance and Accounting* **2**(3), 164–173 (2013)
14. Aksenova, A., van Esch, D., Flynn, J., Golik, P.: How might we create better benchmarks for speech recognition? In: Proceedings of the 1st Workshop on Benchmarking: Past, Present and Future, pp. 22–34 (2021)
15. Akyürek, A.F., Akyürek, E., Madaan, A., Kalyan, A., Clark, P., Wijaya, D., Tandon, N.: RI4f: Generating natural language feedback with reinforcement learning for repairing model outputs (2023)
16. Akyürek, E., Schuurmans, D., Andreas, J., Ma, T., Zhou, D.: What learning algorithm is in-context learning? investigations with linear models. In: The Eleventh International Conference on Learning Representations (2023)

17. Allen-Zhu, Z., Li, Y.: Feature purification: How adversarial training performs robust deep learning. In: 2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS), pp. 977–988. IEEE (2022)
18. Allen-Zhu, Z., Li, Y.: Towards understanding ensemble, knowledge distillation and self-distillation in deep learning. In: The Eleventh International Conference on Learning Representations (2023)
19. Allen-Zhu, Z., Li, Y., Liang, Y.: Learning and generalization in overparameterized neural networks, going beyond two layers. In: Advances in neural information processing systems, pp. 6155–6166 (2019)
20. Allen-Zhu, Z., Li, Y., Song, Z.: A convergence theory for deep learning via over-parameterization. In: International Conference on Machine Learning, pp. 242–252. PMLR (2019)
21. Amari, S.i.: Backpropagation and stochastic gradient descent method. *Neurocomputing* **5**(4–5), 185–196 (1993)
22. An, S., Chou, S.Y., Zhang, K., Xu, Q., Tao, G., Shen, G., Cheng, S., Ma, S., Chen, P.Y., Ho, T.Y., et al.: Elijah: Eliminating backdoors injected in diffusion models via distribution shift. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 38, pp. 10847–10855 (2024)
23. Anderson, M., Anderson, S.L.: Guest editors’ introduction: machine ethics. *IEEE Intelligent Systems* **21**(4), 10–11 (2006)
24. Anderson, M., Anderson, S.L.: Machine ethics: Creating an ethical intelligent agent. *AI magazine* **28**(4), 15–15 (2007)
25. Andrés, M.E., Bordenabe, N.E., Chatzikokolakis, K., Palamidessi, C.: Geoindistinguishability: Differential privacy for location-based systems. In: Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security, pp. 901–914 (2013)
26. Anil, R., Dai, A.M., Firat, O., Johnson, M., Lepikhin, D., Passos, A., Shakeri, S., Taropa, E., Bailey, P., Chen, Z., Chu, E., Clark, J.H., Shafey, L.E., Huang, Y., Meier-Hellstern, K., Mishra, G., Moreira, E., Omernick, M., Robinson, K., Ruder, S., Tay, Y., Xiao, K., Xu, Y., Zhang, Y., Abrego, G.H., Ahn, J., Austin, J., Barham, P., Botha, J., Bradbury, J., Brahma, S., Brooks, K., Catasta, M., Cheng, Y., Cherry, C., Choquette-Choo, C.A., Chowdhery, A., Crepy, C., Dave, S., Dehghani, M., Dev, S., Devlin, J., Díaz, M., Du, N., Dyer, E., Feinberg, V., Feng, F., Fienber, V., Freitag, M., Garcia, X., Gehrmann, S., Gonzalez, L., Gur-Ari, G., Hand, S., Hashemi, H., Hou, L., Howland, J., Hu, A., Hui, J., Hurwitz, J., Isard, M., Ittycheriah, A., Jagielski, M., Jia, W., Kenealy, K., Krikun, M., Kudugunta, S., Lan, C., Lee, K., Lee, B., Li, E., Li, M., Li, W., Li, Y., Li, J., Lim, H., Lin, H., Liu, Z., Liu, F., Maggioni, M., Mahendru, A., Maynez, J., Misra, V., Moussalem, M., Nado, Z., Nham, J., Ni, E., Nystrom, A., Parrish, A., Pellat, M., Polacek, M., Polozov, A., Pope, R., Qiao, S., Reif, E., Richter, B., Riley, P., Ros, A.C., Roy, A., Saeta, B., Samuel, R., Shelby, R., Slone, A., Smilkov, D., So, D.R., Sohn, D., Tokumine, S., Valter, D., Vasudevan, V., Vodrahalli, K., Wang, X., Wang, P., Wang, Z., Wang, T., Wieting, J., Wu, Y., Xu, K., Xu, Y., Xue, L., Yin, P., Yu, J., Zhang, Q., Zheng, S., Zheng, C., Zhou, W., Zhou, D., Petrov, S., Wu, Y.: Palm 2 technical report (2023)
27. Anil, R., Ghazi, B., Gupta, V., Kumar, R., Manurangsi, P.: Large-scale differentially private bert. In: Findings of the Association for Computational Linguistics: EMNLP 2022, pp. 6481–6491 (2022)
28. Appen: Unraveling the link between translations and gender bias in llms (2023). URL <https://appen.com/blog/unraveling-the-link-between-translations-and-gender-bias-in-llms/>
29. Apple, D.: Learning with privacy at scale. *Apple Machine Learning Journal* **1**(8) (2017)
30. Arif, H., Gittens, A., Chen, P.Y.: Reprogrammable-fl: Improving utility-privacy tradeoff in federated learning via model reprogramming. In: 2023 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML), pp. 197–209. IEEE (2023)

31. Arif, H., Gittens, A., Chen, P.Y.: Reprogrammable-FL: Improving utility-privacy tradeoff in federated learning via model reprogramming. In: First IEEE Conference on Secure and Trustworthy Machine Learning (2023)
32. Arisoy, E., Sethy, A., Ramabhadran, B., Chen, S.: Bidirectional recurrent neural network language models for automatic speech recognition. In: 2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 5421–5425. IEEE (2015)
33. Arora, S., Du, S., Hu, W., Li, Z., Wang, R.: Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In: International Conference on Machine Learning, pp. 322–332 (2019)
34. Askell, A., Bai, Y., Chen, A., Drain, D., Ganguli, D., Henighan, T., Jones, A., Joseph, N., Mann, B., DasSarma, N., Elhage, N., Hatfield-Dodds, Z., Hernandez, D., Kernion, J., Ndousse, K., Olsson, C., Amodei, D., Brown, T., Clark, J., McCandlish, S., Olah, C., Kaplan, J.: A general language assistant as a laboratory for alignment (2021)
35. Augenstein, I., Baldwin, T., Cha, M., Chakraborty, T., Ciampaglia, G.L., Corney, D., DiResta, R., Ferrara, E., Hale, S., Halevy, A., et al.: Factuality challenges in the era of large language models. arXiv preprint arXiv:2310.05189 (2023)
36. Bach, F., Jenatton, R., Mairal, J., Obozinski, G., et al.: Optimization with sparsity-inducing penalties. *Foundations and Trends® in Machine Learning* **4**(1), 1–106 (2012)
37. Bahdanau, D., Cho, K., Bengio, Y.: Neural machine translation by jointly learning to align and translate. arXiv preprint arXiv:1409.0473 (2014)
38. Bahng, H., Jahanian, A., Sankaranarayanan, S., Isola, P.: Exploring visual prompts for adapting large-scale models. arXiv preprint arXiv:2203.17274 **1**(3), 4 (2022)
39. Bai, S., Kolter, J.Z., Koltun, V.: An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. arXiv preprint arXiv:1803.01271 (2018)
40. Bai, Y., Chen, F., Wang, H., Xiong, C., Mei, S.: Transformers as statisticians: Provable in-context learning with in-context algorithm selection. arXiv preprint arXiv:2306.04637 (2023)
41. Bai, Y., Jones, A., Ndousse, K., Askell, A., Chen, A., DasSarma, N., Drain, D., Fort, S., Ganguli, D., Henighan, T., Joseph, N., Kadavath, S., Kernion, J., Conerly, T., El-Showk, S., Elhage, N., Hatfield-Dodds, Z., Hernandez, D., Hume, T., Johnston, S., Kravec, S., Lovitt, L., Nanda, N., Olsson, C., Amodei, D., Brown, T., Clark, J., McCandlish, S., Olah, C., Mann, B., Kaplan, J.: Training a helpful and harmless assistant with reinforcement learning from human feedback (2022)
42. Bai, Y., Jones, A., Ndousse, K., Askell, A., Chen, A., DasSarma, N., Drain, D., Fort, S., Ganguli, D., Henighan, T., Joseph, N., Kadavath, S., Kernion, J., Conerly, T., El-Showk, S., Elhage, N., Hatfield-Dodds, Z., Hernandez, D., Hume, T., Johnston, S., Kravec, S., Lovitt, L., Nanda, N., Olsson, C., Amodei, D., Brown, T., Clark, J., McCandlish, S., Olah, C., Mann, B., Kaplan, J.: Training a helpful and harmless assistant with reinforcement learning from human feedback (2022)
43. Bai, Y., Kadavath, S., Kundu, S., Askell, A., Kernion, J., Jones, A., Chen, A., Goldie, A., Mirhoseini, A., McKinnon, C., et al.: Constitutional ai: Harmlessness from ai feedback. arXiv preprint arXiv:2212.08073 (2022)
44. Bai, Y., Zhao, J., Shi, J., Wei, T., Wu, X., He, L.: Fairbench: A four-stage automatic framework for detecting stereotypes and biases in large language models (2023)
45. Baidu: Ernie - baidu yiyan (2023). <https://yiyan.baidu.com/>
46. Balunovic, M., Dimitrov, D.I., Jovanović, N., Vechev, M.: Lamp: Extracting text from gradients with language model priors. In: Advances in Neural Information Processing Systems (2022)
47. Bao, F., Li, C., Zhu, J., Zhang, B.: Analytic-dpm: an analytic estimate of the optimal reverse variance in diffusion probabilistic models. In: ICLR (2022)
48. Baranchuk, D., Voynov, A., Rubachev, I., Khrulkov, V., Babenko, A.: Label-efficient semantic segmentation with diffusion models. In: ICLR (2022)
49. Barbaro, M., Jr., T.Z.: A face is exposed for aol searcher no. 4417749. *New York Times* (2006)

50. Barham, P., Chowdhery, A., Dean, J., Ghemawat, S., Hand, S., Hurt, D., Isard, M., Lim, H., Pang, R., Roy, S., et al.: Pathways: Asynchronous distributed dataflow for ml. *Proceedings of Machine Learning and Systems* **4**, 430–449 (2022)
51. Baydin, A.G., Pearlmutter, B.A., Syme, D., Wood, F., Torr, P.: Gradients without backpropagation. *arXiv preprint arXiv:2202.08587* (2022)
52. Bazaraa, M.S., Sherali, H.D., Shetty, C.M.: *Nonlinear programming: theory and algorithms*. John Wiley & sons (2013)
53. Bedapudi, P.: Nudenet: Neural nets for nudity classification, detection and selective censoring (2019)
54. Bedapudi, P.: Nudenet: Neural nets for nudity classification, detection and selective censoring (2019)
55. Behnia, R., Ebrahimi, M.R., Pacheco, J., Padmanabhan, B.: Ew-tune: A framework for privately fine-tuning large language models with differential privacy. In: *2022 IEEE International Conference on Data Mining Workshops (ICDMW)*, pp. 560–566. IEEE (2022)
56. Belghazi, M.I., Baratin, A., Rajeshwar, S., Ozair, S., Bengio, Y., Courville, A., Hjelm, D.: Mutual information neural estimation. In: *International Conference on Machine Learning*, pp. 531–540 (2018)
57. Berahas, A.S., Cao, L., Choromanski, K., Scheinberg, K.: A theoretical and empirical comparison of gradient approximations in derivative-free optimization. *Found. Comput. Math.* **22**(2), 507–560 (2022)
58. Bevendorff, J., Potthast, M., Hagen, M., Stein, B.: Heuristic authorship obfuscation. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1098–1108 (2019)
59. Bhagoji, A.N., Cullina, D., Mittal, P.: Lower bounds on adversarial robustness from optimal transport. *Advances in Neural Information Processing Systems* **32** (2019)
60. Bhardwaj, R., Poria, S.: Red-teaming large language models using chain of utterances for safety-alignment (2023)
61. Bhattamishra, S., Ahuja, K., Goyal, N.: On the ability and limitations of transformers to recognize formal languages. *arXiv preprint arXiv:2009.11264* (2020)
62. Bhattamishra, S., Patel, A., Goyal, N.: On the computational power of transformers and its implications in sequence modeling. *arXiv preprint arXiv:2006.09286* (2020)
63. Bi, Z., Zhang, N., Xue, Y., Ou, Y., Ji, D., Zheng, G., Chen, H.: Oceangpt: A large language model for ocean science tasks (2023)
64. Bianchi, F., Suzgun, M., Attanasio, G., Röttger, P., Jurafsky, D., Hashimoto, T., Zou, J.: Safety-tuned llamas: Lessons from improving the safety of large language models that follow instructions. *arXiv preprint arXiv:2309.07875* (2023)
65. Blanco-Justicia, A., Sánchez, D., Domingo-Ferrer, J., Muralidhar, K.: A critical review on the use (and misuse) of differential privacy in machine learning. *ACM Computing Surveys* **55**(8), 1–16 (2022)
66. Blier, L., Ollivier, Y.: The description length of deep learning models. *Advances in Neural Information Processing Systems* **31** (2018)
67. Bommasani, R., Hudson, D.A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M.S., Bohg, J., Bosselut, A., Brunskill, E., et al.: On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* (2021)
68. Bommasani, R., Klyman, K., Longpre, S., Kapoor, S., Maslej, N., Xiong, B., Zhang, D., Liang, P.: The foundation model transparency index (2023)
69. Boneh, M.: Vehicle detection using deep learning and yolo algorithm (2023)
70. Borji, A.: A categorical archive of chatgpt failures. *arXiv preprint arXiv:2302.03494* (2023)
71. Bossard, L., Guillaumin, M., Van Gool, L.: Food-101 – mining discriminative components with random forests. In: *European Conference on Computer Vision* (2014)
72. Bourtole, L., Chandrasekaran, V., Choquette-Choo, C.A., Jia, H., Travers, A., Zhang, B., Lie, D., Papernot, N.: Machine unlearning. In: *2021 IEEE Symposium on Security and Privacy (SP)*, pp. 141–159. IEEE (2021)

73. Bovens, M.: Two concepts of accountability: Accountability as a virtue and as a mechanism. *West European Politics* **33**(5), 946–967 (2010)
74. Bowman, S.R., Hyun, J., Perez, E., Chen, E., Pettit, C., Heiner, S., Lukošūtė, K., Askill, A., Jones, A., Chen, A., et al.: Measuring progress on scalable oversight for large language models. *arXiv preprint arXiv:2211.03540* (2022)
75. Box, G.E., Jenkins, G.M., Reinsel, G.C., Ljung, G.M.: *Time series analysis: forecasting and control*. John Wiley & Sons (2015)
76. Boyd, S.P., Vandenberghe, L.: *Convex optimization*. Cambridge university press (2004)
77. Brassil, J., Low, S.H., Maxemchuk, N.F., O’Gorman, L.: Electronic marking and identification techniques to discourage document copying. *IEEE J. Sel. Areas Commun.* **13**(8), 1495–1504 (1995)
78. Brohan, A., Brown, N., Carbajal, J., Chebotar, Y., Chen, X., Choromanski, K., Ding, T., Driess, D., Dubey, A., Finn, C., et al.: Rt-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818* (2023)
79. Brooks, T., Holynski, A., Efros, A.A.: Instructpix2pix: Learning to follow image editing instructions. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18392–18402 (2023)
80. Brown, H., Lee, K., Miresghallah, F., Shokri, R., Tramèr, F.: What does it mean for a language model to preserve privacy? In: *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pp. 2280–2292 (2022)
81. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., Amodei, D.: Language models are few-shot learners. In: *Advances in Neural Information Processing Systems*, vol. 33, pp. 1877–1901 (2020)
82. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. *Advances in Neural Information Processing Systems* **33**, 1877–1901 (2020)
83. Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D.M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., Amodei, D.: Language models are few-shot learners. In: *NeurIPS* (2020)
84. Brutzkus, A., Globerson, A.: An optimization and generalization analysis for max-pooling networks. In: *Uncertainty in Artificial Intelligence*, pp. 1650–1660. PMLR (2021)
85. Buschek, D., Mecke, L., Lehmann, F., Dang, H.: Nine potential pitfalls when designing human-ai co-creative systems. *arXiv preprint arXiv:2104.00358* (2021)
86. Caballero, D., Araya, R., Kronholm, H., Viiri, J., Mansikkaniemi, A., Lehesvuori, S., Virtanen, T., Kurimo, M.: Asr in classroom today: Automatic visualization of conceptual network in science classrooms. In: *Data Driven Approaches in Digital Education: 12th European Conference on Technology Enhanced Learning, EC-TEL 2017, Tallinn, Estonia, September 12–15, 2017, Proceedings 12*, pp. 541–544. Springer (2017)
87. Cao, B., Cao, Y., Lin, L., Chen, J.: Defending against alignment-breaking attacks via robustly aligned llm (2023)
88. Cao, Y., Gu, Q.: Generalization bounds of stochastic gradient descent for wide and deep neural networks. In: *Advances in Neural Information Processing Systems*, pp. 10836–10846 (2019)
89. Cao, Y., Yang, J.: Towards making systems forget with machine unlearning. In: *2015 IEEE symposium on security and privacy*, pp. 463–480. IEEE (2015)
90. Carlini, N., Liu, C., Erlingsson, Ú., Kos, J., Song, D.: The secret sharer: Evaluating and testing unintended memorization in neural networks. In: *28th {USENIX} Security Symposium ({USENIX} Security 19)*, pp. 267–284 (2019)

91. Carlini, N., Nasr, M., Choquette-Choo, C.A., Jagielski, M., Gao, I., Awadalla, A., Koh, P.W., Ippolito, D., Lee, K., Tramer, F., et al.: Are aligned neural networks adversarially aligned? arXiv preprint arXiv:2306.15447 (2023)
92. Carlini, N., Tramer, F., Wallace, E., Jagielski, M., Herbert-Voss, A., Lee, K., Roberts, A., Brown, T., Song, D., Erlingsson, U., et al.: Extracting training data from large language models. In: 30th USENIX Security Symposium (USENIX Security 21), pp. 2633–2650 (2021)
93. Carlucci, F.M., D’Innocente, A., Bucci, S., Caputo, B., Tommasi, T.: Domain generalization by solving jigsaw puzzles. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 2229–2238 (2019)
94. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 9650–9660 (2021)
95. Carranza, A.G., Farahani, R., Ponomareva, N., Kurakin, A., Jagielski, M., Nasr, M.: Privacy-preserving recommender systems with synthetic query generation using differentially private large language models. arXiv preprint arXiv:2305.05973 (2023)
96. Carroll, M., Chan, A., Ashton, H., Krueger, D.: Characterizing manipulation from ai systems. arXiv preprint arXiv:2303.09387 (2023)
97. Carvalho, R.S., Vasiloudis, T., Feyisetan, O.: Tem: high utility metric differential privacy on text. arXiv preprint arXiv:2107.07928 (2021)
98. Casper, S., Lin, J., Kwon, J., Culp, G., Hadfield-Menell, D.: Explore, establish, exploit: Red teaming language models from scratch. arXiv preprint arXiv:2306.09442 (2023)
99. Celikyilmaz, A., Clark, E., Gao, J.: Evaluation of text generation: A survey (2021)
100. Chakraborty, S., Bedi, A.S., Zhu, S., An, B., Manocha, D., Huang, F.: On the possibilities of ai-generated text detection. CoRR **abs/2304.04736** (2023)
101. Challu, C., Olivares, K.G., Oreshkin, B.N., Ramirez, F.G., Canseco, M.M., Dubrawski, A.: Nhits: neural hierarchical interpolation for time series forecasting. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 37, pp. 6989–6997 (2023)
102. Chan, W., Jaitly, N., Le, Q., Vinyals, O.: Listen, attend and spell: A neural network for large vocabulary conversational speech recognition. In: 2016 IEEE international conference on acoustics, speech and signal processing (ICASSP), pp. 4960–4964. IEEE (2016)
103. Chang, C., Peng, W.C., Chen, T.F.: Llm4ts: Two-stage fine-tuning for time-series forecasting with pre-trained llms. arXiv preprint arXiv:2308.08469 (2023)
104. Chang, Y., Wang, X., Wang, J., Wu, Y., Yang, L., Zhu, K., Chen, H., Yi, X., Wang, C., Wang, Y., Ye, W., Zhang, Y., Chang, Y., Yu, P.S., Yang, Q., Xie, X.: A survey on evaluation of large language models (2023)
105. Chao, P., Robey, A., Dobriban, E., Hassani, H., Pappas, G.J., Wong, E.: Jailbreaking black box large language models in twenty queries. CoRR **abs/2310.08419** (2023)
106. Chatzikokolakis, K., Andrés, M.E., Bordenabe, N.E., Palamidessi, C.: Broadening the scope of differential privacy using metrics. In: Privacy Enhancing Technologies: 13th International Symposium, PETS 2013, Bloomington, IN, USA, July 10–12, 2013. Proceedings 13, pp. 82–102. Springer (2013)
107. Chen, A., Lorenz, P., Yao, Y., Chen, P.Y., Liu, S.: Visual prompting for adversarial robustness. In: ICASSP 2023–2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 1–5. IEEE (2023)
108. Chen, A., Yao, Y., Chen, P.Y., Zhang, Y., Liu, S.: Understanding and improving visual prompting: A label-mapping perspective. arXiv preprint arXiv:2211.11635 (2022)
109. Chen, A., Zhang, Y., Jia, J., Diffenderfer, J., Liu, J., Parasyris, K., Zhang, Y., Zhang, Z., Kailkhura, B., Liu, S.: Deepzero: Scaling up zeroth-order optimization for deep model training. ICLR (2024)
110. Chen, A., Zhang, Y., Jia, J., Diffenderfer, J., Parasyris, K., Liu, J., Zhang, Y., Zhang, Z., Kailkhura, B., Liu, S.: Deepzero: Scaling up zeroth-order optimization for deep model training. In: The Twelfth International Conference on Learning Representations (2024)

111. Chen, C., Feng, X., Zhou, J., Yin, J., Zheng, X.: Federated large language model: A position paper. arXiv preprint arXiv:2307.08925 (2023)
112. Chen, C., Hou, N., Hu, Y., Shirol, S., Chng, E.S.: Noise-robust speech recognition with 10 minutes unparallelized in-domain data. In: ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 4298–4302. IEEE (2022)
113. Chen, C., Hu, Y., Yang, C.H.H., Siniscalchi, S.M., Chen, P.Y., Chng, E.: Hyporadise: An open baseline for generative speech recognition with large language models. In: Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track (2023)
114. Chen, C., Shu, K.: Combating misinformation in the age of llms: Opportunities and challenges. arXiv preprint arXiv:2311.05656 (2023)
115. Chen, H., Lu, C., Ying, C., Su, H., Zhu, J.: Offline reinforcement learning via high-fidelity generative behavior modeling. In: ArXiv (2022)
116. Chen, L., Lu, K., Rajeswaran, A., Lee, K., Grover, A., Laskin, M., Abbeel, P., Srinivas, A., Mordatch, I.: Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems* **34**, 15084–15097 (2021)
117. Chen, L.C., Papandreou, G., Schroff, F., Adam, H.: Rethinking atrous convolution for semantic image segmentation. arXiv preprint arXiv:1706.05587 (2017)
118. Chen, M., Gao, W., Liu, G., Peng, K., Wang, C.: Boundary unlearning: Rapid forgetting of deep networks via shifting the decision boundary. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 7766–7775 (2023)
119. Chen, M., Radford, A., Child, R., Wu, J., Jun, H., Luan, D., Sutskever, I.: Generative pretraining from pixels. In: H.D. III, A. Singh (eds.) *Proceedings of the 37th International Conference on Machine Learning, Proceedings of Machine Learning Research*, vol. 119, pp. 1691–1703. PMLR (2020). URL <https://proceedings.mlr.press/v119/chen20s.html>
120. Chen, P.Y.: Model reprogramming: Resource-efficient cross-domain machine learning **38**(20), 22584–22591 (2024)
121. Chen, P.Y., Das, P.: AI Maintenance: A robustness perspective. *Computer* **56**(2), 48–56 (2023)
122. Chen, P.Y., Hsieh, C.J.: *Adversarial Robustness for Machine Learning*. Academic Press (2022)
123. Chen, P.Y., Liu, S.: Holistic adversarial robustness of deep learning models. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 37, pp. 15411–15420 (2023)
124. Chen, P.Y., Zhang, H., Sharma, Y., Yi, J., Hsieh, C.J.: ZOO: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In: ACM Workshop on Artificial Intelligence and Security, pp. 15–26 (2017)
125. Chen, Q., Zhao, H., Li, W., Huang, P., Ou, W.: Behavior sequence transformer for e-commerce recommendation in alibaba. In: Proceedings of the 1st International Workshop on Deep Learning Practice for High-Dimensional Sparse Data, pp. 1–4 (2019)
126. Chen, S., Mo, F., Wang, Y., Chen, C., Nie, J.Y., Wang, C., Cui, J.: A customized text sanitization mechanism with differential privacy. In: Findings of the Association for Computational Linguistics: ACL 2023, pp. 5747–5758 (2023)
127. Chen, S., Sun, P., Song, Y., Luo, P.: Diffusiondet: Diffusion model for object detection. In: ArXiv (2022)
128. Chen, S., Wang, C., Chen, Z., Wu, Y., Liu, S., Chen, Z., Li, J., Kanda, N., Yoshioka, T., Xiao, X., et al.: Wavlm: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing* **16**(6), 1505–1518 (2022)
129. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International Conference on Machine Learning (2018)
130. Chen, T., Kornblith, S., Swersky, K., Norouzi, M., Hinton, G.E.: Big self-supervised models are strong semi-supervised learners. In: *NeruIPS* (2020)
131. Chen, T., Zhang, Z., Zhang, Y., Chang, S., Liu, S., Wang, Z.: Quarantine: Sparsity can uncover the trojan attack trigger for free. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 598–609 (2022)

132. Chen, W., Song, D., Li, B.: Trojdiff: Trojan attacks on diffusion models with diverse targets. In: CVPR (2023)
133. Chen, X., Hsieh, C.J., Gong, B.: When vision transformers outperform resnets without pre-training or strong data augmentations. In: International Conference on Learning Representations (2021)
134. Chen, X., Liu, S., Xu, K., Li, X., Lin, X., Hong, M., Cox, D.: Zo-adamm: Zeroth-order adaptive momentum method for black-box optimization. In: Advances in Neural Information Processing Systems, pp. 7202–7213 (2019)
135. Chen, Y., Wang, Z., Xing, X., huimin zheng, Xu, Z., Fang, K., Wang, J., Li, S., Wu, J., Liu, Q., Xu, X.: Bianque: Balancing the questioning and suggestion ability of health llms with multi-turn health conversations polished by chatgpt (2023)
136. Chen, Z., Cao, Y., Gu, Q., Zhang, T.: A generalized neural tangent kernel analysis for two-layer neural networks. *Advances in Neural Information Processing Systems* **33** (2020)
137. Chen, Z.C., Yang, C.H.H., Li, B., Zhang, Y., Chen, N., Chang, S.Y., Prabhavalkar, R., yi Lee, H., Sainath, T.: How to Estimate Model Transferability of Pre-Trained Speech Models? In: Proc. INTERSPEECH 2023, pp. 456–460 (2023). <https://doi.org/10.21437/Interspeech.2023-1079>
138. Cheng, M., Le, T., Chen, P., Zhang, H., Yi, J., Hsieh, C.: Query-efficient hard-label black-box attack: An optimization-based approach. In: 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019. OpenReview.net (2019)
139. Cheng, X., Chen, Y., Sra, S.: Transformers implement functional gradient descent to learn non-linear functions in context. arXiv preprint arXiv:2312.06528 (2023)
140. Chi, C., Feng, S., Du, Y., Xu, Z., Cousineau, E., Burchfiel, B., Song, S.: Diffusion policy: Visuomotor policy learning via action diffusion (2023)
141. Chin, Z.Y., Jiang, C.M., Huang, C.C., Chen, P.Y., Chiu, W.C.: Prompting4debugging: Red-teaming text-to-image diffusion models by finding problematic prompts. arXiv preprint arXiv:2309.06135 (2023)
142. Chin, Z.Y., Jiang, C.M., Huang, C.C., Chen, P.Y., Chiu, W.C.: Prompting4debugging: Red-teaming text-to-image diffusion models by finding problematic prompts. arXiv preprint arXiv:2309.06135 (2023)
143. Chin, Z.Y., Jiang, C.M., Huang, C.C., Chen, P.Y., Chiu, W.C.: Prompting4debugging: Red-teaming text-to-image diffusion models by finding problematic prompts. *International Conference on Machine Learning* (2024)
144. Chorowski, J., Jaitly, N.: Towards better decoding and language model integration in sequence to sequence models. arXiv preprint arXiv:1612.02695 (2016)
145. Chou, S.Y., Chen, P.Y., Ho, T.Y.: How to backdoor diffusion models? In: CVPR (2023)
146. Chou, S.Y., Chen, P.Y., Ho, T.Y.: Villandiffusion: A unified backdoor attack framework for diffusion models. *Advances in Neural Information Processing Systems* **36** (2023)
147. Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., Barham, P., Chung, H.W., Sutton, C., Gehrmann, S., Schuh, P., Shi, K., Tsvyashchenko, S., Maynez, J., Rao, A., Barnes, P., Tay, Y., Shazeer, N., Prabhakaran, V., Reif, E., Du, N., Hutchinson, B., Pope, R., Bradbury, J., Austin, J., Isard, M., Gur-Ari, G., Yin, P., Duke, T., Levskaya, A., Ghemawat, S., Dev, S., Michalewski, H., Garcia, X., Misra, V., Robinson, K., Fedus, L., Zhou, D., Ippolito, D., Luan, D., Lim, H., Zoph, B., Spiridonov, A., Sepassi, R., Dohan, D., Agrawal, S., Omernick, M., Dai, A.M., Pillai, T.S., Pellat, M., Lewkowycz, A., Moreira, E., Child, R., Polozov, O., Lee, K., Zhou, Z., Wang, X., Saeta, B., Diaz, M., Firat, O., Catasta, M., Wei, J., Meier-Hellstern, K., Eck, D., Dean, J., Petrov, S., Fiedel, N.: Palm: Scaling language modeling with pathways (2022)
148. Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., Barham, P., Chung, H.W., Sutton, C., Gehrmann, S., et al.: Palm: Scaling language modeling with pathways. arXiv preprint arXiv:2204.02311 (2022)

149. Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., Barham, P., Chung, H.W., Sutton, C., Gehrmann, S., et al.: Palm: Scaling language modeling with pathways. arXiv preprint arXiv:2204.02311 (2022)
150. Christ, M., Gunn, S., Zamir, O.: Undetectable watermarks for language models. arXiv preprint arXiv:2306.09194 (2023)
151. Christie, G., Fendley, N., Wilson, J., Mukherjee, R.: Functional map of the world. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2018)
152. Chrupala, G.: Putting natural in natural language processing. arXiv preprint arXiv:2305.04572 (2023)
153. Chu, X., Tian, Z., Wang, Y., Zhang, B., Ren, H., Wei, X., Xia, H., Shen, C.: Twins: Revisiting the design of spatial attention in vision transformers. *Advances in Neural Information Processing Systems* **34**, 9355–9366 (2021)
154. Chu, Z., Hao, H., Ouyang, X., Wang, S., Wang, Y., Shen, Y., Gu, J., Cui, Q., Li, L., Xue, S., et al.: Leveraging large language models for pre-trained recommender systems. arXiv preprint arXiv:2308.10837 (2023)
155. Chung, H.W., Hou, L., Longpre, S., Zoph, B., Tay, Y., Fedus, W., Li, Y., Wang, X., Dehghani, M., Brahma, S., et al.: Scaling instruction-finetuned language models. arXiv preprint arXiv:2210.11416 (2022)
156. Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., , Vedaldi, A.: Describing textures in the wild. In: Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (2014)
157. Clusmann, J., Kolbinger, F.R., Muti, H.S., Carrero, Z.I., Eckardt, J.N., Laleh, N.G., Löffler, C.M.L., Schwarzkopf, S.C., Unger, M., Veldhuizen, G.P., et al.: The future landscape of large language models in medicine. *Communications Medicine* **3**(1), 141 (2023)
158. CMU: Enron email dataset (2015). <https://www.cs.cmu.edu/~enron/>
159. Codella, N., Rotemberg, V., Tschandl, P., Celebi, M.E., Dusza, S., Gutman, D., Helba, B., Kalloo, A., Liopyris, K., Marchetti, M., et al.: Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (isic). arXiv preprint arXiv:1902.03368 (2019)
160. Cohen, J., Rosenfeld, E., Kolter, J.Z.: Certified adversarial robustness via randomized smoothing. In: K. Chaudhuri, R. Salakhutdinov (eds.) Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9–15 June 2019, Long Beach, California, USA, *Proceedings of Machine Learning Research*, vol. 97, pp. 1310–1320. PMLR (2019)
161. Colon-Hernandez, P., Lieberman, H., Xin, Y., Yin, C., Breazeal, C., Chin, P.: Adversarial transformer language models for contextual commonsense inference. *CoRR* **abs/2302.05406** (2023)
162. Conover, M., Hayes, M., Mathur, A., Xie, J., Wan, J., Shah, S., Ghodsi, A., Wendell, P., Zaharia, M., Xin, R.: Free dolly: Introducing the world's first truly open instruction-tuned llm (2023). URL <https://www.databricks.com/blog/2023/04/12/dolly-first-open-commercially-viable-instruction-tuned-llm>
163. Cooper, A.F., Moss, E., Laufer, B., Nissenbaum, H.: Accountability in an algorithmic society: relationality, responsibility, and robustness in machine learning. In: Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency, pp. 864–876 (2022)
164. Cordonnier, J.B., Loukas, A., Jaggi, M.: On the relationship between self-attention and convolutional layers. In: International Conference on Learning Representations (2019)
165. Council, F.T.: 10 ways cybercriminals can abuse large language models (2023). URL <https://www.forbes.com/sites/forbestechcouncil/2023/06/30/10-ways-cybercriminals-can-abuse-large-language-models/>
166. Council, F.T.: Navigating the biases in llm generative ai: A guide to responsible implementation (2023). URL <https://www.forbes.com/sites/forbestechcouncil/2023/09/06/navigating-the-biases-in-llm-generative-ai-a-guide-to-responsible-implementation/>

167. Cui, L., Wu, Y., Liu, J., Yang, S., Zhang, Y.: Template-based named entity recognition using bart. In: Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021, pp. 1835–1845 (2021)
168. Dahlmeier, D., Ng, H.T.: Better evaluation for grammatical error correction. In: Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pp. 568–572 (2012)
169. Dai, L., Mao, J., Fan, X., Zhou, X.: Deepfider: A multi-module and invisibility watermarking scheme for language model. *CoRR* **abs/2208.04676** (2022)
170. Dai, W., Li, J., Li, D., Tiong, A.M.H., Zhao, J., Wang, W., Li, B., Fung, P., Hoi, S.: Instructblip: Towards general-purpose vision-language models with instruction tuning (2023)
171. Dai, W., Lin, J., Jin, F., Li, T., Tsai, Y.S., Gasevic, D., Chen, G.: Can large language models provide feedback to students? a case study on chatgpt (2023). <https://doi.org/10.35542/osf.io/hcgzj>
172. Dan, C., Wei, Y., Ravikumar, P.: Sharp statistical guarantees for adversarially robust gaussian classification. In: International Conference on Machine Learning, pp. 2345–2355. PMLR (2020)
173. Daniely, A., Malach, E.: Learning parities with neural networks. *Advances in Neural Information Processing Systems* **33**, 20356–20365 (2020)
174. Dasgupta, I., Lampinen, A.K., Chan, S.C.Y., Sheahan, H.R., Creswell, A., Kumaran, D., McClelland, J.L., Hill, F.: Language models show human-like content effects on reasoning tasks (2023)
175. De Laat, P.B.: Algorithmic decision-making based on machine learning from big data: can transparency restore accountability? *Philosophy & technology* **31**(4), 525–541 (2018)
176. Dehghani, M., Gouws, S., Vinyals, O., Uszkoreit, J., Kaiser, L.: Universal transformers. In: International Conference on Learning Representations (2018)
177. Deldari, S., Xue, H., Saeed, A., He, J., Smith, D.V., Salim, F.D.: Beyond just vision: A review on self-supervised representation learning on multimodal and temporal data. *arXiv preprint arXiv:2206.02353* (2022)
178. Deng, J., Wang, Y., Li, J., Wang, C., Shang, C., Liu, H., Rajasekaran, S., Ding, C.: Tag: Gradient attack on transformer-based language models. In: Findings of the Association for Computational Linguistics: EMNLP 2021, pp. 3600–3610 (2021)
179. Deshpande, A., Murahari, V., Rajpurohit, T., Kalyan, A., Narasimhan, K.: Toxicity in chatgpt: Analyzing persona-assigned language models. *arXiv preprint arXiv:2304.05335* (2023)
180. Dettmers, T., Pagnoni, A., Holtzman, A., Zettlemoyer, L.: Qlora: Efficient finetuning of quantized llms. *arXiv preprint arXiv:2305.14314* (2023)
181. Dev, S., Jha, A., Goyal, J., Tewari, D., Dave, S., Prabhakaran, V.: Building stereotype repositories with llms and community engagement for scale and depth. *Cross-Cultural Considerations in NLP@ EACL* p. 84 (2023)
182. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018)
183. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding (2019)
184. Dhariwal, P., Nichol, A.Q.: Diffusion models beat gans on image synthesis. In: NIPS (2021)
185. Dhingra, H., Jayashanker, P., Moghe, S., Strubell, E.: Queer people are people first: Deconstructing sexual identity stereotypes in large language models. *arXiv preprint arXiv:2307.00101* (2023)
186. Dhurandhar, A., Pedapati, T., Balakrishnan, A., Chen, P.Y., Shanmugam, K., Puri, R.: Model agnostic contrastive explanations for structured data. *arXiv preprint arXiv:1906.00117* (2019)
187. Dinh, L., Sohl-Dickstein, J., Bengio, S.: Density estimation using real NVP. In: ICLR (2017)
188. Dobriban, E., Hassani, H., Hong, D., Robey, A.: Provable tradeoffs in adversarially robust classification. *arXiv preprint arXiv:2006.05161* (2020)

189. Dolatabadi, H.M., Erfani, S.M., Leckie, C.: Advflow: Inconspicuous black-box adversarial attacks using normalizing flows. In: NIPS (2020)
190. Dong, H., Xiong, W., Goyal, D., Pan, R., Diao, S., Zhang, J., Shum, K., Zhang, T.: Raft: Reward ranked finetuning for generative foundation model alignment. arXiv preprint arXiv:2304.06767 (2023)
191. Dong, L., Xu, S., Xu, B.: Speech-transformer: a no-recurrence sequence-to-sequence model for speech recognition. In: 2018 IEEE international conference on acoustics, speech and signal processing (ICASSP), pp. 5884–5888. IEEE (2018)
192. Dong, Y., Cordonnier, J.B., Loukas, A.: Attention is not all you need: Pure attention loses rank doubly exponentially with depth. In: International Conference on Machine Learning, pp. 2793–2803. PMLR (2021)
193. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
194. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2020)
195. Du, S.S., Zhai, X., Poczos, B., Singh, A.: Gradient descent provably optimizes over-parameterized neural networks. In: International Conference on Learning Representations (2019). URL <https://openreview.net/forum?id=S1eK3i09YQ>
196. Du, Y., Li, S., Torralba, A., Tenenbaum, J.B., Mordatch, I.: Improving factuality and reasoning in language models through multiagent debate. arXiv preprint arXiv:2305.14325 (2023)
197. Dubois, Y., Hashimoto, T., Ermon, S., Liang, P.: Improving self-supervised learning by characterizing idealized representations (2022). <https://doi.org/10.48550/ARXIV.2209.06235>. URL <https://arxiv.org/abs/2209.06235>
198. Dubois, Y., Kiela, D., Schwab, D.J., Vedantam, R.: Learning optimal representations with the decodable information bottleneck. Advances in Neural Information Processing Systems **33**, 18674–18690 (2020)
199. Duchi, J.C., Jordan, M.I., Wainwright, M.J.: Local privacy and statistical minimax rates. In: 2013 IEEE 54th Annual Symposium on Foundations of Computer Science, pp. 429–438. IEEE (2013)
200. Duchi, J.C., Jordan, M.I., Wainwright, M.J., Wibisono, A.: Optimal rates for zero-order convex optimization: The power of two function evaluations. IEEE Transactions on Information Theory **61**(5), 2788–2806 (2015)
201. Dutta, S., Jain, S., Maheshwari, A., Pal, S., Ramakrishnan, G., Jyothi, P.: Error correction in asr using sequence-to-sequence models. arXiv preprint arXiv:2202.01157 (2022)
202. Dwork, C., McSherry, F., Nissim, K., Smith, A.: Calibrating noise to sensitivity in private data analysis. In: Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4–7, 2006. Proceedings 3, pp. 265–284. Springer (2006)
203. Dwork, C., Roth, A., et al.: The algorithmic foundations of differential privacy. Foundations and Trends® in Theoretical Computer Science **9**(3–4), 211–407 (2014)
204. Edelman, B.L., Goel, S., Kakade, S., Zhang, C.: Inductive biases and variable creation in self-attention mechanisms. In: International Conference on Machine Learning, pp. 5793–5831. PMLR (2022)
205. Eldan, R., Russinovich, M.: Who’s harry potter? approximate unlearning in llms (2023)
206. Elsayed, G.F., Goodfellow, I., Sohl-Dickstein, J.: Adversarial reprogramming of neural networks. In: International Conference on Learning Representations (2019)
207. Erlingsson, Ú., Pihur, V., Korolova, A.: Rappor: Randomized aggregatable privacy-preserving ordinal response. In: Proceedings of the 2014 ACM SIGSAC conference on computer and communications security, pp. 1054–1067 (2014)
208. Ethayarajh, K.: How contextual are contextualized word representations? comparing the geometry of bert, elmo, and GPT-2 embeddings. In: K. Inui, J. Jiang, V. Ng, X. Wan

- (eds.) Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019, pp. 55–65. Association for Computational Linguistics (2019). <https://doi.org/10.18653/v1/D19-1006>
209. Face, H.: The big benchmarks collection - a open-llm-leaderboard collection. <https://huggingface.co/spaces/OpenLLMBenchmark/The-Big-Benchmarks-Collection>
 210. Fan, A., Lewis, M., Dauphin, Y.N.: Hierarchical neural story generation. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers, pp. 889–898 (2018)
 211. Fan, A., Lewis, M., Dauphin, Y.N.: Hierarchical neural story generation. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers, pp. 889–898 (2018)
 212. Fan, A., Lewis, M., Dauphin, Y.N.: Hierarchical neural story generation. In: I. Gurevych, Y. Miyao (eds.) Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers, pp. 889–898. Association for Computational Linguistics (2018). <https://aclanthology.org/P18-1082/>
 213. Fan, C., Liu, J., Hero, A., Liu, S.: Challenging forgets: Unveiling the worst-case forget sets in machine unlearning. European Conference on Computer Vision (ECCV) (2024)
 214. Fan, C., Liu, J., Zhang, Y., Wei, D., Wong, E., Liu, S.: Salun: Empowering machine unlearning via gradient-based weight saliency in both image classification and generation. In: International Conference on Learning Representations (2024)
 215. Fan, W., Zhao, Z., Li, J., Liu, Y., Mei, X., Wang, Y., Wen, Z., Wang, F., Zhao, X., Tang, J., Li, Q.: Recommender systems in the era of large language models (llms) (2023)
 216. Fang, Y., Wang, W., Xie, B., Sun, Q., Wu, L., Wang, X., Huang, T., Wang, X., Cao, Y.: Eva: Exploring the limits of masked visual representation learning at scale. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 19358–19369 (2023)
 217. Fawaz, H.I., Forestier, G., Weber, J., Idoumghar, L., Muller, P.A.: Transfer learning for time series classification. In: 2018 IEEE international conference on big data (Big Data), pp. 1367–1376. IEEE (2018)
 218. Fei, Z., Shen, X., Zhu, D., Zhou, F., Han, Z., Zhang, S., Chen, K., Shen, Z., Ge, J.: Lawbench: Benchmarking legal knowledge of large language models. arXiv preprint arXiv:2309.16289 (2023)
 219. Feldman, P., Dant, A., Rosenbluth, D.: Ethics, rules of engagement, and ai: Neural narrative mapping using large transformer language models. arXiv preprint arXiv:2202.02647 (2022)
 220. Feldman, V., Zhang, C.: What neural networks memorize and why: Discovering the long tail via influence estimation. Advances in Neural Information Processing Systems **33**, 2881–2891 (2020)
 221. Felkner, V.K., Chang, H.C.H., Jang, E., May, J.: Winoqueer: A community-in-the-loop benchmark for anti-lgbtq+ bias in large language models. arXiv preprint arXiv:2306.15087 (2023)
 222. Fermi, E.: Numerical solution of a minimum problem. Tech. rep., Los Alamos Scientific Lab., Los Alamos, NM (1952)
 223. Fernandez, P., Chaffin, A., Tit, K., Chappelier, V., Furon, T.: Three bricks to consolidate watermarks for large language models (2023)
 224. Feyisetan, O., Balle, B., Drake, T., Diethe, T.: Privacy-and utility-preserving textual analysis via calibrated multivariate perturbations. In: Proceedings of the 13th International Conference on Web Search and Data Mining, pp. 178–186 (2020)
 225. Fisher, R.A.: On the interpretation of χ^2 from contingency tables, and the calculation of p. Journal of the royal statistical society **85**(1), 87–94 (1922)
 226. Fjeld, J., Achten, N., Hilligoss, H., Nagy, Á., Srikumar, M.: Principled artificial intelligence: Mapping consensus in ethical and rights-based approaches to principles for ai. SSRN Electronic Journal (2020)

227. Flaxman, A.D., Kalai, A.T., McMahan, H.B.: Online convex optimization in the bandit setting: gradient descent without a gradient. In: ACM-SIAM symposium on Discrete algorithms, pp. 385–394. Society for Industrial and Applied Mathematics (2005)
228. Foret, P., Kleiner, A., Mobahi, H., Neyshabur, B.: Sharpness-aware minimization for efficiently improving generalization. In: International Conference on Learning Representations (2020)
229. Frank, M.: Baby steps in evaluating the capacities of large language models. *Nature Reviews Psychology* **2** (2023). <https://doi.org/10.1038/s44159-023-00211-x>
230. Frankle, J., Carbin, M.: The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv preprint arXiv:1803.03635* (2018)
231. Frankle, J., Dziugaite, G.K., Roy, D., Carbin, M.: Linear mode connectivity and the lottery ticket hypothesis. In: International Conference on Machine Learning, pp. 3259–3269. PMLR (2020)
232. Frantar, E., Alistarh, D.: Sparsegpt: Massive language models can be accurately pruned in one-shot. In: International Conference on Machine Learning, pp. 10323–10337. PMLR (2023)
233. Frieder, S., Pinchetti, L., Chevalier, A., Griffiths, R.R., Salvatori, T., Lukasiewicz, T., Petersen, P.C., Berner, J.: Mathematical capabilities of chatgpt (2023)
234. Fu, H., Chi, Y., Liang, Y.: Guaranteed recovery of one-hidden-layer neural networks via cross entropy. *IEEE Transactions on Signal Processing* **68**, 3225–3235 (2020)
235. Fu, L., Chai, H., Luo, S., Du, K., Zhang, W., Fan, L., Lei, J., Rui, R., Lin, J., Fang, Y., Liu, Y., Wang, J., Qi, S., Zhang, K., Zhang, W., Yu, Y.: Codeapex: A bilingual programming evaluation benchmark for large language models (2023)
236. Fu, S.W., Liao, C.F., Tsao, Y., Lin, S.D.: Metricgan: Generative adversarial networks based black-box metric scores optimization for speech enhancement. In: International Conference on Machine Learning, pp. 2031–2041. PMLR (2019)
237. Fu, Y., Peng, H., Khot, T., Lapata, M.: Improving language model negotiation with self-play and in-context learning from ai feedback. *arXiv preprint arXiv:2305.10142* (2023)
238. Gal, R., Alaluf, Y., Atzmon, Y., Patashnik, O., Bermano, A.H., Chechik, G., Cohen-or, D.: An image is worth one word: Personalizing text-to-image generation using textual inversion. In: ICLR (2022)
239. Gallegos, I.O., Rossi, R.A., Barrow, J., Tanjim, M.M., Kim, S., Derroncourt, F., Yu, T., Zhang, R., Ahmed, N.K.: Bias and fairness in large language models: A survey (2023)
240. Gandhe, A., Rastrow, A.: Audio-attention discriminative language model for asr rescoring. In: ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 7944–7948. IEEE (2020)
241. Gandikota, R., Materzynska, J., Fiotto-Kaufman, J., Bau, D.: Erasing concepts from diffusion models. In: CVPR, pp. 2426–2436 (2023)
242. Gandikota, R., Materzynska, J., Fiotto-Kaufman, J., Bau, D.: Erasing concepts from diffusion models. In: International Conference on Computer Vision (ICCV) (2023)
243. Gandikota, R., Orgad, H., Belinkov, Y., Materzyńska, J., Bau, D.: Unified concept editing in diffusion models. *arXiv preprint arXiv:2308.14761* (2023)
244. Ganguli, D., Lovitt, L., Kernion, J., Askell, A., Bai, Y., Kadavath, S., Mann, B., Perez, E., Schiefer, N., Ndousse, K., et al.: Red teaming language models to reduce harms: Methods, scaling behaviors, and lessons learned. *arXiv preprint arXiv:2209.07858* (2022)
245. Gao, L., Tow, J., Abbasi, B., Biderman, S., Black, S., DiPofi, A., Foster, C., Golding, L., Hsu, J., Le Noac’h, A., Li, H., McDonnell, K., Muennighoff, N., Ociepa, C., Phang, J., Reynolds, L., Schoelkopf, H., Skowron, A., Sutow, L., Tang, E., Thite, A., Wang, B., Wang, K., Zou, A.: A framework for few-shot language model evaluation (2023). <https://zenodo.org/records/10256836>
246. Gao, P., Han, J., Zhang, R., Lin, Z., Geng, S., Zhou, A., Zhang, W., Lu, P., He, C., Yue, X., et al.: Llama-adapter v2: Parameter-efficient visual instruction model. *arXiv preprint arXiv:2304.15010* (2023)

247. Gao, T., Fisch, A., Chen, D.: Making pre-trained language models better few-shot learners. In: Association for Computational Linguistics (ACL) (2021)
248. Garg, S., Tsipras, D., Liang, P.S., Valiant, G.: What can transformers learn in-context? a case study of simple function classes. *Advances in Neural Information Processing Systems* **35**, 30583–30598 (2022)
249. Gehman, S., Gururangan, S., Sap, M., Choi, Y., Smith, N.A.: Realtotoxicityprompts: Evaluating neural toxic degeneration in language models. *arXiv preprint arXiv:2009.11462* (2020)
250. Gehman, S., Gururangan, S., Sap, M., Choi, Y., Smith, N.A.: Realtotoxicityprompts: Evaluating neural toxic degeneration in language models. *arXiv preprint arXiv:2009.11462* (2020)
251. Gehrmann, S., Adewumi, T., Aggarwal, K., Ammanamanchi, P.S., Anuoluwapo, A., Bosse-lut, A., Chandu, K.R., Clinciu, M., Das, D., Dhole, K.D., Du, W., Durmus, E., Dušek, O., Emezue, C., Gangal, V., Garbacea, C., Hashimoto, T., Hou, Y., Jernite, Y., Jhamtani, H., Ji, Y., Jolly, S., Kale, M., Kumar, D., Ladhak, F., Madaan, A., Maddela, M., Mahajan, K., Mahamood, S., Majumder, B.P., Martins, P.H., McMillan-Major, A., Mille, S., van Miltenburg, E., Nadeem, M., Narayan, S., Nikolaev, V., Niyongabo, R.A., Osei, S., Parikh, A., Perez-Beltrachini, L., Rao, N.R., Raunak, V., Rodriguez, J.D., Santhanam, S., Sedoc, J., Sellam, T., Shaikh, S., Shimorina, A., Cabezudo, M.A.S., Strobelt, H., Subramani, N., Xu, W., Yang, D., Yerukola, A., Zhou, J.: The gem benchmark: Natural language generation, its evaluation and metrics (2021)
252. Gehrmann, S., Strobelt, H., Rush, A.M.: GLTR: statistical detection and visualization of generated text. In: *ACL*, pp. 111–116 (2019)
253. Geiping, J., Garrido, Q., Fernandez, P., Bar, A., Pirsiavash, H., LeCun, Y., Goldblum, M.: A cookbook of self-supervised learning. *arXiv preprint arXiv:2304.12210* (2023)
254. Ghadimi, S., Lan, G.: Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization* **23**(4), 2341–2368 (2013)
255. Gilad-Bachrach, R., Burges, C.J.: The median hypothesis. Tech. Rep. MSR-TR-2012-56 (2012). URL <https://www.microsoft.com/en-us/research/publication/the-median-hypothesis/>
256. Gilson, A., Safranek, C., Huang, T., Socrates, V., Chi, L., Taylor, R., Chartash, D.: How does chatgpt perform on the united states medical licensing examination? the implications of large language models for medical education and knowledge assessment. *JMIR medical education* **9**, e45312 (2023). <https://doi.org/10.2196/45312>
257. Goel, S., Prabhu, A., Sanyal, A., Lim, S.N., Torr, P., Kumaraguru, P.: Towards adversarial evaluations for inexact machine unlearning. *arXiv preprint arXiv:2201.06640* (2022)
258. Gokaslan, A., Cohen, V., Pavlick, E., Tellex, S.: Openwebtext corpus. <http://Skyilion007.github.io/OpenWebTextCorpus> (2019)
259. Golatkar, A., Achille, A., Soatto, S.: Eternal sunshine of the spotless net: Selective forgetting in deep networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9304–9312 (2020)
260. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. *Communications of the ACM* **63**(11), 139–144 (2020)
261. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. *International Conference on Learning Representations* (2015)
262. Gould, S., Fernando, B., Cherian, A., Anderson, P., Cruz, R.S., Guo, E.: On differentiating parameterized argmin and argmax problems with application to bi-level optimization. *arXiv preprint arXiv:1607.05447* (2016)
263. Graves, A.: Sequence transduction with recurrent neural networks. *arXiv preprint arXiv:1211.3711* (2012)
264. Graves, A., Fernández, S., Gomez, F., Schmidhuber, J.: Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In: *Proceedings of the 23rd international conference on Machine learning*, pp. 369–376 (2006)
265. Greenslade, T.A., Félix-Brasdefer, J.C.: Error correction and learner perceptions in 12 spanish writing. In: *Selected Proceedings of the 7th Conference on the Acquisition of*

- Spanish and Portuguese as First and Second Languages, pp. 185–194. Somerville, MA: Cascadilla Proceedings Project (2006)
266. Grefenstette, J.J.: Genetic algorithms and machine learning. In: Proceedings of the sixth annual conference on Computational learning theory, pp. 3–4 (1993)
 267. Greshake, K., Abdelnabi, S., Mishra, S., Endres, C., Holz, T., Fritz, M.: Not what you’ve signed up for: Compromising real-world llm-integrated applications with indirect prompt injection. In: Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security, pp. 79–90 (2023)
 268. Grosse, R., Bae, J., Anil, C., Elhage, N., Tamkin, A., Tajdini, A., Steiner, B., Li, D., Durmus, E., Perez, E., et al.: Studying large language model generalization with influence functions. arXiv preprint arXiv:2308.03296 (2023)
 269. Gruver, N., Finzi, M.A., Qiu, S., Wilson, A.G.: Large language models are zero-shot time series forecasters. Advances in Neural Information Processing Systems (2023)
 270. Gu, B., Liu, G., Zhang, Y., Geng, X., Huang, H.: Optimizing large-scale hyperparameters via automated learning algorithm. arXiv preprint arXiv:2102.09026 (2021)
 271. Gu, J., Feng, C., Zhao, Z., Ying, Z., Chen, R.T., Pan, D.Z.: Efficient on-chip learning for optical neural networks through power-aware sparse zeroth-order optimization. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 35, pp. 7583–7591 (2021)
 272. Gu, Z., Zhu, X., Ye, H., Zhang, L., Wang, J., Jiang, S., Xiong, Z., Li, Z., He, Q., Xu, R., Huang, W., Wang, Z., Wang, S., Zheng, W., Feng, H., Xiao, Y.: Xiezhi: An ever-updating benchmark for holistic domain knowledge evaluation (2023)
 273. Guha, N., Nyarko, J., Ho, D.E., Ré, C., Chilton, A., Narayana, A., Chohlas-Wood, A., Peters, A., Waldon, B., Rockmore, D.N., Zambrano, D., Talisman, D., Hoque, E., Surani, F., Fagan, F., Sarfaty, G., Dickinson, G.M., Porat, H., Hegland, J., Wu, J., Nudell, J., Niklaus, J., Nay, J., Choi, J.H., Tobia, K., Hagan, M., Ma, M., Livermore, M., Rasumov-Rahe, N., Holzenberger, N., Kolt, N., Henderson, P., Rehaag, S., Goel, S., Gao, S., Williams, S., Gandhi, S., Zur, T., Iyer, V., Li, Z.: Legalbench: A collaboratively built benchmark for measuring legal reasoning in large language models (2023)
 274. Gulati, A., Qin, J., Chiu, C.C., Parmar, N., Zhang, Y., Yu, J., Han, W., Wang, S., Zhang, Z., Wu, Y., et al.: Conformer: Convolution-augmented transformer for speech recognition. Proc. Interspeech 2020 pp. 5036–5040 (2020)
 275. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: International conference on machine learning, pp. 1321–1330. PMLR (2017)
 276. Guo, E., Gupta, M., Sinha, S., Rössler, K., Tatagiba, M., Akagami, R., Al-Mefty, O., Sugiyama, T., Stieg, P.E., Pickett, G.E., et al.: neurogpt-x: Towards an accountable expert opinion tool for vestibular schwannoma. medRxiv pp. 2023–02 (2023)
 277. Guo, T., Guo, K., Nan, B., Liang, Z., Guo, Z., Chawla, N.V., Wiest, O., Zhang, X.: What can large language models do in chemistry? a comprehensive benchmark on eight tasks. In: NeurIPS (2023)
 278. Guo, T., Hu, W., Mei, S., Wang, H., Xiong, C., Savarese, S., Bai, Y.: How do transformers learn in-context beyond simple functions? a case study on learning with representations. arXiv preprint arXiv:2310.10616 (2023)
 279. Guo, Z., Wei, Y., Liu, M., Ji, Z., Bai, J., Guo, Y., Zuo, W.: Black-box tuning of vision-language models with effective gradient approximation. In: EMNLP, pp. 5356–5368 (2023)
 280. Gupta, K., Jawalkar, N., Mukherjee, A., Chandran, N., Gupta, D., Panwar, A., Sharma, R.: Sigma: Secure gpt inference with function secret sharing. Cryptology ePrint Archive, Paper 2023/1269 (2023). URL <https://eprint.iacr.org/2023/1269>. <https://eprint.iacr.org/2023/1269>
 281. Hadfield-Menell, D., Russell, S.J., Abbeel, P., Dragan, A.: Cooperative inverse reinforcement learning. Advances in neural information processing systems **29** (2016)
 282. Han, S., Mao, H., Dally, W.J.: Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. arXiv preprint arXiv:1510.00149 (2015)
 283. Han, S., Mao, H., Dally, W.J.: Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. arXiv preprint arXiv:1510.00149 (2015)

284. HaoChen, J.Z., Ma, T.: A theoretical study of inductive biases in contrastive learning. arXiv preprint arXiv:2211.14699 (2022)
285. Hartvigsen, T., Gabriel, S., Palangi, H., Sap, M., Ray, D., Kamar, E.: Toxigen: A large-scale machine-generated dataset for adversarial and implicit hate speech detection. arXiv preprint arXiv:2203.09509 (2022)
286. Hassan, H., Aue, A., Chen, C., Chowdhary, V., Clark, J., Federmann, C., Huang, X., Junczys-Dowmunt, M., Lewis, W., Li, M., et al.: Achieving human parity on automatic chinese to english news translation. arXiv preprint arXiv:1803.05567 (2018)
287. Hayes, M.H.: Statistical digital signal processing and modeling. John Wiley & Sons (1996)
288. Hazell, J.: Large language models can be used to effectively scale spear phishing campaigns. arXiv preprint arXiv:2305.06972 (2023)
289. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 16000–16009 (2022)
290. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 770–778 (2016)
291. Heckerman, D.: Bayesian networks for data mining. Data mining and knowledge discovery **1**, 79–119 (1997)
292. Helber, P., Bischke, B., Dengel, A., Borth, D.: Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing (2019)
293. Hendrycks, D., Burns, C., Basart, S., Zou, A., Mazeika, M., Song, D., Steinhardt, J.: Measuring massive multitask language understanding. arXiv preprint arXiv:2009.03300 (2020)
294. Hendrycks, D., Gimpel, K.: Gaussian error linear units (gelus). arXiv preprint arXiv:1606.08415 (2016)
295. Hessel, J., Holtzman, A., Forbes, M., Le Bras, R., Choi, Y.: Clipscore: A reference-free evaluation metric for image captioning. In: EMNLP, pp. 7514–7528 (2021)
296. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: NIPS (2017)
297. Hlubík, P., Španěl, M., Boháč, M., Weingartová, L.: Inserting punctuation to asr output in a real-time production environment. In: Text, Speech, and Dialogue: 23rd International Conference, TSD 2020, Brno, Czech Republic, September 8–11, 2020, Proceedings, pp. 418–425. Springer (2020)
298. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in neural information processing systems **33**, 6840–6851 (2020)
299. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: NIPS (2020)
300. Ho, J., Saharia, C., Chan, W., Fleet, D.J., Norouzi, M., Salimans, T.: Cascaded diffusion models for high fidelity image generation. In: JMLR (2022)
301. Ho, J., Salimans, T.: Classifier-free diffusion guidance. In: NIPS Workshop on Deep Generative Models and Downstream Applications (2021)
302. Ho, J., Salimans, T.: Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598 (2022)
303. Ho, J., Salimans, T., Gritsenko, A.A., Chan, W., Norouzi, M., Fleet, D.J.: Video diffusion models. In: NeurIPS (2022)
304. Hoang, D.T., Chollampatt, S., Ng, H.T.: Exploiting n-best hypotheses to improve an smt approach to grammatical error correction. arXiv preprint arXiv:1606.00210 (2016)
305. Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural computation **9**(8), 1735–1780 (1997)
306. Hoffmann, J., Borgeaud, S., Mensch, A., Buchatskaya, E., Cai, T., Rutherford, E., Casas, D.d.L., Hendricks, L.A., Welbl, J., Clark, A., et al.: Training compute-optimal large language models. arXiv preprint arXiv:2203.15556 (2022)

307. Holmes, J., Liu, Z., Zhang, L., Ding, Y., Sio, T.T., McGee, L.A., Ashman, J.B., Li, X., Liu, T., Shen, J., Liu, W.: Evaluating large language models on a highly-specialized topic, radiation oncology physics. *Frontiers in Oncology* **13** (2023). <https://doi.org/10.3389/fonc.2023.1219326>
308. Holtzman, A., Buys, J., Du, L., Forbes, M., Choi, Y.: The curious case of neural text degeneration. In: 8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020 (2020)
309. Holtzman, A., Buys, J., Du, L., Forbes, M., Choi, Y.: The curious case of neural text degeneration. In: 8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020 (2020)
310. Hoofnagle, C.J., van der Sloot, B., Borgesius, F.Z.: The european union general data protection regulation: what it is and what it means. *Information & Communications Technology Law* **28**(1), 65–98 (2019)
311. Hou, A.B., Zhang, J., He, T., Wang, Y., Chuang, Y.S., Wang, H., Shen, L., Durme, B.V., Khashabi, D., Tsvetkov, Y.: Semstamp: A semantic watermark with paraphrastic robustness for text generation (2023)
312. Hou, X., Liu, J., Li, J., Li, Y., Jie Lu, W., Hong, C., Ren, K.: Ciphergpt: Secure two-party gpt inference. *Cryptology ePrint Archive*, Paper 2023/1147 (2023). URL <https://eprint.iacr.org/2023/1147>
313. Houben, S., Stallkamp, J., Salmen, J., Schlipsing, M., Igel, C.: Detection of traffic signs in real-world images: The german traffic sign detection benchmark. In: The 2013 International Joint Conference on Neural Networks (IJCNN), pp. 1–8 (2013)
314. Hounsby, N., Giurgiu, A., Jastrzebski, S., Morrone, B., De Laroussilhe, Q., Gesmundo, A., Attariyan, M., Gelly, S.: Parameter-efficient transfer learning for nlp. In: International Conference on Machine Learning, pp. 2790–2799. PMLR (2019)
315. Hovy, D., Johannsen, A., Søgaard, A.: User review sites as a resource for large-scale sociolinguistic studies. In: Proceedings of the 24th international conference on World Wide Web, pp. 452–461 (2015)
316. Howard, J.: Imagenette: A smaller subset of 10 easily classified classes from imagenet (2019)
317. Howard, J., Ruder, S.: Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146* (2018)
318. Hron, J., Bahri, Y., Sohl-Dickstein, J., Novak, R.: Infinite attention: Nngp and ntk for deep attention networks. In: International Conference on Machine Learning, pp. 4376–4386. PMLR (2020)
319. Hsiung, L., Tang, Y.C., Chen, P.Y., Ho, T.Y.: NCTV: Neural Clamping Toolkit and Visualization for Neural Network Calibration. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 37 (13), pp. 16446–16448 (2023)
320. Hsu, C.Y., Tsai, Y.L., Lin, C.H., Chen, P.Y., Yu, C.M., Huang, C.Y.: Safe lora: the silver lining of reducing safety risks when fine-tuning large language models. *arXiv preprint arXiv:2405.16833* (2024)
321. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Chen, W.: Lora: Low-rank adaptation of large language models (2021). URL <https://arxiv.org/abs/2106.09685>
322. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685* (2021)
323. Hu, K., Pang, R., Sainath, T.N., Strohmaier, T.: Transformer based deliberation for two-pass speech recognition. In: 2021 IEEE Spoken Language Technology Workshop (SLT), pp. 68–74. IEEE (2021)
324. Hu, K., Sainath, T.N., Pang, R., Prabhavalkar, R.: Deliberation model based two-pass end-to-end speech recognition. In: ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 7799–7803. IEEE (2020)
325. Hu, K., Zou, A., Wang, Z., Leino, K., Fredrikson, M.: Scaling in depth: Unlocking robustness certification on imagenet. *Advances in Neural Information Processing Systems* (2023)

326. Hu, X., Chen, J., Li, X., Guo, Y., Wen, L., Yu, P.S., Guo, Z.: Do large language models know about facts? arXiv preprint arXiv:2310.05177 (2023)
327. Hu, X., Chen, P.Y., Ho, T.Y.: Radar: Robust ai-text detection via adversarial learning. *Advances in Neural Information Processing Systems* **36**, 15077–15095 (2023)
328. Hu, X., Chen, P.Y., Ho, T.Y.: Gradient cuff: Detecting jailbreak attacks on large language models by exploring refusal loss landscapes. arXiv preprint arXiv:2403.00867 (2024)
329. Hu, Y., Chen, C., Li, R., Zhu, Q., Chng, E.S.: Noise-aware speech enhancement using diffusion probabilistic model. arXiv preprint arXiv:2307.08029 (2023)
330. Hu, Y., Chen, C., Yang, C.H.H., Li, R., Zhang, C., Chen, P.Y., Chng, E.: Large language models are efficient learners of noise-robust speech recognition. *International Conference on Learning Representations* (2024)
331. Hu, Y., Chen, C., Zhu, Q., Chng, E.S.: Wav2code: Restore clean speech representations via codebook lookup for noise-robust asr. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* (2023)
332. Hu, Y., Hou, N., Chen, C., Chng, E.S.: Interactive feature fusion for end-to-end noise-robust speech recognition. In: *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6292–6296. IEEE (2022)
333. Hu, Y., Hou, N., Chen, C., Chng, E.S.: Dual-Path Style Learning for End-to-End Noise-Robust Speech Recognition. In: *Proc. INTERSPEECH 2023*, pp. 2918–2922 (2023)
334. Huang, H., Peng, F.: An empirical study of efficient asr rescoring with transformers. arXiv preprint arXiv:1910.11450 (2019)
335. Huang, J., Chang, K.C.C.: Citation: A key to building responsible and accountable large language models. arXiv preprint arXiv:2307.02185 (2023)
336. Huang, J., Shao, H., Chang, K.C.C.: Are large pre-trained language models leaking your personal information? (2022)
337. Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., Chen, Q., Peng, W., Feng, X., Qin, B., et al.: A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. arXiv preprint arXiv:2311.05232 (2023)
338. Huang, R., Lam, M.W.Y., Wang, J., Su, D., Yu, D., Ren, Y., Zhao, Z.: Fastdiff: A fast conditional diffusion model for high-quality speech synthesis. In: *IJCAI* (2022)
339. Huang, S., Jiang, Z., Dong, H., Qiao, Y., Gao, P., Li, H.: Instruct2act: Mapping multi-modality instructions to robotic actions with large language model. arXiv preprint arXiv:2305.11176 (2023)
340. Huang, Y., Bai, Y., Zhu, Z., Zhang, J., Zhang, J., Su, T., Liu, J., Lv, C., Zhang, Y., Lei, J., Fu, Y., Sun, M., He, J.: C-eval: A multi-level multi-discipline chinese evaluation suite for foundation models (2023)
341. Huang, Y., Cheng, Y., Liang, Y.: In-context convergence of transformers. In: *NeurIPS 2023 Workshop on Mathematics of Modern Machine Learning* (2023)
342. Huang, Y., Gupta, S., Xia, M., Li, K., Chen, D.: Catastrophic jailbreak of open-source llms via exploiting generation. arXiv preprint arXiv:2310.06987 (2023)
343. Huang, Y., Gupta, S., Xia, M., Li, K., Chen, D.: Catastrophic jailbreak of open-source llms via exploiting generation (2023)
344. Huang, Y., Shi, J., Li, Y., Fan, C., Wu, S., Zhang, Q., Liu, Y., Zhou, P., Wan, Y., Gong, N.Z., et al.: Metatool benchmark for large language models: Deciding whether to use tools and which to use. arXiv preprint arXiv:2310.03128 (2023)
345. Huang, Y., Su, Y., Ravi, S., Song, Z., Arora, S., Li, K.: Privacy-preserving learning via deep net pruning. arXiv preprint arXiv:2003.01876 (2020)
346. Igamberdiev, T., Habernal, I.: DP-BART for privatized text rewriting under local differential privacy. In: *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 13914–13934. Association for Computational Linguistics, Toronto, Canada (2023)
347. Ilharco, G., Ribeiro, M.T., Wortsman, M., Gururangan, S., Schmidt, L., Hajishirzi, H., Farhadi, A.: Editing models with task arithmetic. arXiv preprint arXiv:2212.04089 (2022)

348. Ilharco, G., Wortsman, M., Wightman, R., Gordon, C., Carlini, N., Taori, R., Dave, A., Shankar, V., Namkoong, H., Miller, J., Hajishirzi, H., Farhadi, A., Schmidt, L.: Openclip (2021). <https://doi.org/10.5281/zenodo.5143773>
349. Ilyas, A., Engstrom, L., Athalye, A., Lin, J.: Black-box adversarial attacks with limited queries and information. International Conference on International Conference on Machine Learning (2018)
350. Ilyas, A., Engstrom, L., Athalye, A., Lin, J.: Black-box adversarial attacks with limited queries and information. In: J.G. Dy, A. Krause (eds.) Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10–15, 2018, *Proceedings of Machine Learning Research*, vol. 80, pp. 2142–2151. PMLR (2018)
351. Ioffe, S.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. arXiv preprint arXiv:1502.03167 (2015)
352. Ippolito, D., Duckworth, D., Callison-Burch, C., Eck, D.: Automatic detection of generated text is easiest when humans are fooled. In: ACL, pp. 1808–1822 (2020)
353. Islam, P., Kannappan, A., Kiela, D., Qian, R., Scherrer, N., Vidgen, B.: Financebench: A new benchmark for financial question answering (2023)
354. Izzo, Z., Smart, M.A., Chaudhuri, K., Zou, J.: Approximate data deletion from machine learning models. In: International Conference on Artificial Intelligence and Statistics, pp. 2008–2016. PMLR (2021)
355. Jacot, A., Gabriel, F., Hongler, C.: Neural tangent kernel: Convergence and generalization in neural networks. In: Advances in neural information processing systems, pp. 8571–8580 (2018)
356. Jain, N., Schwarzschild, A., Wen, Y., Somepalli, G., Kirchenbauer, J., Chiang, P., Goldblum, M., Saha, A., Geiping, J., Goldstein, T.: Baseline defenses for adversarial attacks against aligned language models. CoRR **abs/2309.00614** (2023)
357. Jalil, S., Rafi, S., LaToza, T.D., Moran, K., Lam, W.: Chatgpt and software testing education: Promises & perils. In: 2023 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW), pp. 4130–4137. IEEE (2023)
358. James Vuckovic Baratin Aristide, R.T.d.C.: A mathematical theory of attention. arXiv preprint arXiv:2007.02876 (2020)
359. Jang, E., Gu, S., Poole, B.: Categorical reparameterization with gumbel-softmax. In: ICLR (2017)
360. Jang, J., Yoon, D., Yang, S., Cha, S., Lee, M., Logeswaran, L., Seo, M.: Knowledge unlearning for mitigating privacy risks in language models. arXiv preprint arXiv:2210.01504 (2022)
361. Janner, M., Du, Y., Tenenbaum, J.B., Levine, S.: Planning with diffusion for flexible behavior synthesis. In: ICML (2022)
362. Janner, M., Li, Q., Levine, S.: Reinforcement learning as one big sequence modeling problem. In: ICML 2021 Workshop on Unsupervised Reinforcement Learning (2021)
363. Jayaraman, B., Evans, D.: Evaluating differentially private machine learning in practice. In: USENIX Security Symposium (2019)
364. Jelassi, S., Sander, M., Li, Y.: Vision transformers provably learn spatial structure. Advances in Neural Information Processing Systems **35**, 37822–37836 (2022)
365. Jeong, M., Kim, H., Cheon, S.J., Choi, B.J., Kim, N.S.: Diff-tts: A denoising diffusion model for text-to-speech. In: ISCA (2021)
366. Ji, J., Liu, M., Dai, J., Pan, X., Zhang, C., Bian, C., Sun, R., Wang, Y., Yang, Y.: Beavertails: Towards improved safety alignment of llm via a human-preference dataset. arXiv preprint arXiv:2307.04657 (2023)
367. Ji, J., Qiu, T., Chen, B., Zhang, B., Lou, H., Wang, K., Duan, Y., He, Z., Zhou, J., Zhang, Z., Zeng, F., Ng, K.Y., Dai, J., Pan, X., O’Gara, A., Lei, Y., Xu, H., Tse, B., Fu, J., McAleer, S., Yang, Y., Wang, Y., Zhu, S.C., Guo, Y., Gao, W.: Ai alignment: A comprehensive survey (2023)

368. Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y.J., Madotto, A., Fung, P.: Survey of hallucination in natural language generation. *ACM Computing Surveys* **55**(12), 1–38 (2023)
369. Jia, J., Liu, J., Ram, P., Yao, Y., Liu, G., Liu, Y., Sharma, P., Liu, S.: Model sparsity can simplify machine unlearning. In: *Thirty-seventh Conference on Neural Information Processing Systems* (2023)
370. Jia, J., Zhang, Y., Zhang, Y., Liu, J., Runwal, B., Diffenderfer, J., Kailkhura, B., Liu, S.: Soul: Unlocking the power of second-order optimization for llm unlearning. *arXiv preprint arXiv:2404.18239* (2024)
371. Jia, M., Tang, L., Chen, B.C., Cardie, C., Belongie, S., Hariharan, B., Lim, S.N.: Visual prompt tuning. In: *European Conference on Computer Vision*, pp. 709–727. Springer (2022)
372. Jiang, A.Q., Sablayrolles, A., Mensch, A., Bamford, C., Chaplot, D.S., de las Casas, D., Bressand, F., Lengyel, G., Lample, G., Saulnier, L., Lavaud, L.R., Lachaux, M.A., Stock, P., Scao, T.L., Lavril, T., Wang, T., Lacroix, T., Sayed, W.E.: Mistral 7b (2023)
373. Jiang, A.Q., Sablayrolles, A., Mensch, A., Bamford, C., Chaplot, D.S., de las Casas, D., Bressand, F., Lengyel, G., Lample, G., Saulnier, L., Lavaud, L.R., Lachaux, M.A., Stock, P., Scao, T.L., Lavril, T., Wang, T., Lacroix, T., Sayed, W.E.: Mistral 7b (2023)
374. Jiang, Y., Huang, Z., Pan, X., Loy, C.C., Liu, Z.: Talk-to-edit: Fine-grained facial editing via dialog. In: *ICCV* (2021)
375. Jiang, Z., Xu, F.F., Araki, J., Neubig, G.: How can we know what language models know? *Transactions of the Association for Computational Linguistics* **8**, 423–438 (2020)
376. Jiang, Z.H., Hou, Q., Yuan, L., Zhou, D., Shi, Y., Jin, X., Wang, A., Feng, J.: All tokens matter: Token labeling for training better vision transformers. *Advances in Neural Information Processing Systems* **34**, 18590–18602 (2021)
377. Jin, D., Jin, Z., Zhou, J.T., Szolovits, P.: Is bert really robust? a strong baseline for natural language attack on text classification and entailment. In: *Proceedings of the AAAI conference on artificial intelligence*, vol. 34, pp. 8018–8025 (2020)
378. Jin, H., Wei, W., Wang, X., Zhang, W., Wu, Y.: Rethinking learning rate tuning in the era of large language models. *arXiv preprint arXiv:2309.08859* (2023)
379. Jin, M., Wang, S., Ma, L., Chu, Z., Zhang, J.Y., Shi, X., Chen, P.Y., Liang, Y., Li, Y.F., Pan, S., Wen, Q.: Time-LLM: Time series forecasting by reprogramming large language models. In: *The Twelfth International Conference on Learning Representations* (2024)
380. Johnson, R.A., Wichern, D.W., et al.: *Applied multivariate statistical analysis*, vol. 5. Prentice hall Upper Saddle River, NJ (2002)
381. Jones, R., Kumar, R., Pang, B., Tomkins, A.: I know what you did last summer: query logs and user privacy. In: *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*, pp. 909–914 (2007)
382. Joshi, M., Choi, E., Weld, D.S., Zettlemoyer, L.: Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. *arXiv preprint arXiv:1705.03551* (2017)
383. Kankanhalli, M.S., Hau, K.F.: Watermarking of electronic text documents. *Electron. Commer. Res.* **2**(1-2), 169–187 (2002)
384. Kannan, A., Wu, Y., Nguyen, P., Sainath, T.N., Chen, Z., Prabhavalkar, R.: An analysis of incorporating an external language model into a sequence-to-sequence model. In: *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5828. IEEE (2018)
385. Kaplan, J., McCandlish, S., Henighan, T., Brown, T.B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J., Amodei, D.: Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361* (2020)
386. Karp, S., Winston, E., Li, Y., Singh, A.: Local signal adaptivity: Provable feature learning in neural networks beyond kernels. *Advances in Neural Information Processing Systems* **34**, 24883–24897 (2021)
387. Karras, T., Aittala, M., Aila, T., Laine, S.: Elucidating the design space of diffusion-based generative models. In: *NIPS* (2022)

388. Kasiviswanathan, S.P., Lee, H.K., Nissim, K., Raskhodnikova, S., Smith, A.: What can we learn privately? *SIAM Journal on Computing* **40**(3), 793–826 (2011)
389. Katz, G., Barrett, C., Dill, D.L., Julian, K., Kochenderfer, M.J.: Reluplex: An efficient smt solver for verifying deep neural networks. In: *International Conference on Computer Aided Verification*, pp. 97–117. Springer (2017)
390. Kenton, J.D.M.W.C., Toutanova, L.K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: *Proceedings of NAACL-HLT*, pp. 4171–4186 (2019)
391. Kerrigan, G., Slack, D., Tuyls, J.: Differentially private language models benefit from public pre-training. In: *Proceedings of the Second Workshop on Privacy in NLP*, pp. 39–45 (2020)
392. Kešelj, V., Peng, F., Cercone, N., Thomas, C.: N-gram-based author profiles for authorship attribution. In: *Proceedings of the conference pacific association for computational linguistics*, *PACLING*, vol. 3, pp. 255–264 (2003)
393. Khashabi, D., Chaturvedi, S., Roth, M., Upadhyay, S., Roth, D.: Looking beyond the surface: a challenge set for reading comprehension over multiple sentences. In: *Proceedings of North American Chapter of the Association for Computational Linguistics (NAACL)* (2018)
394. Khowaja, S.A., Khuwaja, P., Dev, K.: Chatgpt needs spade (sustainability, privacy, digital divide, and ethics) evaluation: A review. *arXiv preprint arXiv:2305.03123* (2023)
395. Khurana, D., Koli, A., Khatter, K., Singh, S.: Natural language processing: State of the art, current trends and challenges. *Multimedia tools and applications* **82**(3), 3713–3744 (2023)
396. Kim, B., Cai, H., McKenzie, D., Yin, W.: Curvature-aware derivative-free optimization. *arXiv preprint arXiv:2109.13391* (2021)
397. Kim, H., Kim, S., Yoon, S.: Guided-tts: A diffusion model for text-to-speech via classifier guidance. In: *ICML* (2022)
398. Kim, H., Papamakarios, G., Mnih, A.: The lipschitz constant of self-attention. In: *International Conference on Machine Learning*, pp. 5562–5571. PMLR (2021)
399. Kim, J.K., Chua, M., Rickard, M., Lorenzo, A.: Chatgpt and large language model (llm) chatbots: the current state of acceptability and a proposal for guidelines on utilization in academic medicine. *Journal of Pediatric Urology* (2023)
400. Kim, S., Yun, S., Lee, H., Gubri, M., Yoon, S., Oh, S.J.: Propile: Probing privacy leakage in large language models (2023)
401. Kim, S., Yun, S., Lee, H., Gubri, M., Yoon, S., Oh, S.J.: Propile: Probing privacy leakage in large language models. *arXiv preprint arXiv:2307.01881* (2023)
402. Kim, T., Kim, J., Tae, Y., Park, C., Choi, J.H., Choo, J.: Reversible instance normalization for accurate time-series forecasting against distribution shift. In: *International Conference on Learning Representations* (2021)
403. King, M.: Meet dan — the ‘jailbreak’ version of chatgpt and how to use it — ai unchained and unfiltered. <https://medium.com/@neonforge/meet-dan-the-jailbreak-version-of-chatgpt-and-how-to-use-it-ai-unchained-and-unfiltered-f91bfa679024> (2023)
404. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. *International Conference on Learning Representations* (2015)
405. Kingma, D.P., Dhariwal, P.: Glow: Generative flow with invertible 1x1 convolutions. In: *NIPS* (2018)
406. Kingma, D.P., Salimans, T., Poole, B., Ho, J.: Variational diffusion models (2021)
407. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013)
408. Kirchenbauer, J., Geiping, J., Wen, Y., Katz, J., Miers, I., Goldstein, T.: A watermark for large language models. *arXiv preprint arXiv:2301.10226* (2023)
409. Kirchenbauer, J., Geiping, J., Wen, Y., Katz, J., Miers, I., Goldstein, T.: A watermark for large language models. *CoRR abs/2301.10226* (2023)
410. Kirchenbauer, J., Geiping, J., Wen, Y., Shu, M., Saifullah, K., Kong, K., Fernando, K., Saha, A., Goldblum, M., Goldstein, T.: On the reliability of watermarks for large language models (2023)

411. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al.: Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences* **114**(13), 3521–3526 (2017)
412. Kitaev, N., Kaiser, L., Levskaya, A.: Reformer: The efficient transformer. In: *International Conference on Learning Representations* (2020)
413. Ko, C.Y., Chen, P.Y., Das, P., Chuang, Y.S., Daniel, L.: On robustness-accuracy characterization of large language models using synthetic datasets (2023)
414. Ko, C.Y., Chen, P.Y., Das, P., Mohapatra, J., Daniel, L.: What would gauss say about representations? probing pretrained image models using synthetic gaussian benchmarks. In: *International Conference on Machine Learning* (2024)
415. Kocmi, T., Federmann, C.: Large language models are state-of-the-art evaluators of translation quality. In: *Proceedings of the 24th Annual Conference of the European Association for Machine Translation, EAMT 2023, Tampere, Finland, 12-15 June 2023*, pp. 193–203. European Association for Machine Translation (2023). URL <https://aclanthology.org/2023.eamt-1.19>
416. Koh, P.W., Liang, P.: Understanding black-box predictions via influence functions. *International Conference on Machine Learning* (2018)
417. Kojima, T., Gu, S.S., Reid, M., Matsuo, Y., Iwasawa, Y.: Large language models are zero-shot reasoners. In: *Advances in Neural Information Processing Systems* (2022)
418. Kong, Z., Ping, W., Huang, J., Zhao, K., Catanzaro, B.: Diffwave: A versatile diffusion model for audio synthesis. In: *ICLR* (2021)
419. Köpf, A., Kilcher, Y., (ontocord), H.N., Schuhmann, C.: an open assistant for everyone by laion (2023). <https://open-assistant.io/>
420. Krantz, S.G., Parks, H.R.: *The implicit function theorem: history, theory, and applications*. Springer Science & Business Media (2002)
421. Krishna, G., Tran, C., Yu, J., Tewfik, A.H.: Speech recognition with no speech or with noisy speech. In: *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1090–1094. IEEE (2019)
422. Krishna, K., Song, Y., Karpinska, M., Wieting, J., Iyyer, M.: Paraphrasing evades detectors of ai-generated text, but retrieval is an effective defense. *CoRR* **abs/2303.13408** (2023)
423. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images. Tech. rep., University of Toronto, Toronto, Ontario (2009)
424. Krizhevsky, A., Nair, V., Hinton, G.: *Cifar-10* (canadian institute for advanced research) URL <http://www.cs.toronto.edu/~kriz/cifar.html>
425. Kuditipudi, R., Thickstun, J., Hashimoto, T., Liang, P.: Robust distortion-free watermarks for language models. *arXiv preprint arXiv:2307.15593* (2023)
426. Kumar, A., Agarwal, C., Srinivas, S., Feizi, S., Lakkaraju, H.: Certifying LLM safety against adversarial prompting. *CoRR* **abs/2309.02705** (2023)
427. Kumari, N., Zhang, B., Wang, S.Y., Shechtman, E., Zhang, R., Zhu, J.Y.: Ablating concepts in text-to-image diffusion models. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 22691–22702 (2023)
428. Kumari, N., Zhang, B., Wang, S.Y., Shechtman, E., Zhang, R., Zhu, J.Y.: Ablating concepts in text-to-image diffusion models. In: *International Conference on Computer Vision (ICCV)* (2023)
429. Kung, T.H., Cheatham, M., Medenilla, A., Sillos, C., De Leon, L., Elepaño, C., Madriaga, M., Aggabao, R., Diaz-Candido, G., Maningo, J., Tseng, V.: Performance of chatgpt on usmle: Potential for ai-assisted medical education using large language models. *PLOS Digital Health* **2**(2), 1–12 (2023). <https://doi.org/10.1371/journal.pdig.0000198>
430. Kurmanji, M., Triantafillou, P., Triantafillou, E.: Towards unbounded machine unlearning. *arXiv preprint arXiv:2302.09880* (2023)
431. Kwiatkowski, T., Palomaki, J., Redfield, O., Collins, M., Parikh, A., Alberti, C., Epstein, D., Polosukhin, I., Devlin, J., Lee, K., et al.: Natural questions: a benchmark for question

- answering research. *Transactions of the Association for Computational Linguistics* **7**, 453–466 (2019)
432. Lahat, A., Shachar, E., Avidan, B., Shatz, Z., Glicksberg, B., Klang, E.: Evaluating the use of large language model in identifying top research questions in gastroenterology. *Scientific Reports* **13** (2023). <https://doi.org/10.1038/s41598-023-31412-2>
 433. Langer, M., Oster, D., Speith, T., Hermanns, H., Kästner, L., Schmidt, E., Sasing, A., Baum, K.: What do we want from explainable artificial intelligence (xai)?—a stakeholder perspective on xai and a conceptual model guiding interdisciplinary xai research. *Artificial Intelligence* **296**, 103473 (2021)
 434. Lanzi, P.L., Loiacono, D.: Chatgpt and other large language models as evolutionary engines for online interactive collaborative game design (2023)
 435. Lapid, R., Langberg, R., Sipper, M.: Open sesame! universal black box jailbreaking of large language models. *arXiv preprint arXiv:2309.01446* (2023)
 436. Laskar, M.T.R., Bari, M.S., Rahman, M., Bhuiyan, M.A.H., Joty, S., Huang, J.X.: A systematic study and comprehensive evaluation of chatgpt on benchmark datasets (2023)
 437. Latif, S., Qadir, J., Qayyum, A., Usama, M., Younis, S.: Speech technology for healthcare: Opportunities, challenges, and state of the art. *IEEE Reviews in Biomedical Engineering* **14**, 342–356 (2020)
 438. Lavergne, T., Urvoy, T., Yvon, F.: Detecting fake content with relative entropy scoring. In: *Proceedings of the ECAI'08 Workshop on Uncovering Plagiarism, Authorship and Social Software Misuse*, Patras, Greece, July 22, 2008, *CEUR Workshop Proceedings*, vol. 377 (2008)
 439. Le, V.H., Zhang, H.: Log parsing: How far can chatgpt go? (2023)
 440. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**(7553), 436–444 (2015)
 441. Lee, D., Lee, J., Ha, J.W., Kim, J.H., Lee, S.W., Lee, H., Song, H.O.: Query-efficient black-box red teaming via bayesian optimization. In: *Annual Meeting of the Association for Computational Linguistics (ACL)* (2023)
 442. Lee, H., Phatale, S., Mansoor, H., Lu, K., Mesnard, T., Bishop, C., Carbune, V., Rastogi, A.: Rlaif: Scaling reinforcement learning from human feedback with ai feedback. *arXiv preprint arXiv:2309.00267* (2023)
 443. Lee, T., Hong, S., Ahn, J., Hong, I., Lee, H., Yun, S., Shin, J., Kim, G.: Who wrote this code? watermarking for code generation. *CoRR abs/2305.15060* (2023)
 444. Lee, T., Hong, S., Ahn, J., Hong, I., Lee, H., Yun, S., Shin, J., Kim, G.: Who wrote this code? watermarking for code generation (2023)
 445. Lei, Y., Chen, J., Li, S.E., Zheng, S.: Zeroth-order actor-critic. *arXiv preprint arXiv:2201.12518* (2022)
 446. Lei, Y., Lian, J., Yao, J., Huang, X., Lian, D., Xie, X.: Recexplainer: Aligning large language models for recommendation model interpretability (2023)
 447. Leike, J., Sutskever, I.: Introducing Superalignment. <https://openai.com/blog/introducing-superalignment> (2023)
 448. Leng, Y., Tan, X., Liu, W., Song, K., Wang, R., Li, X.Y., Qin, T., Lin, E., Liu, T.Y.: Softcorrect: Error correction with soft detection for automatic speech recognition. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, pp. 13034–13042 (2023)
 449. Leng, Y., Tan, X., Wang, R., Zhu, L., Xu, J., Liu, W., Liu, L., Qin, T., Li, X.Y., Lin, E., et al.: Fastcorrect 2: Fast error correction on multiple candidates for automatic speech recognition. *arXiv preprint arXiv:2109.14420* (2021)
 450. Leng, Y., Tan, X., Zhu, L., Xu, J., Luo, R., Liu, L., Qin, T., Li, X., Lin, E., Liu, T.Y.: Fastcorrect: Fast error correction with edit alignment for automatic speech recognition. *Advances in Neural Information Processing Systems* **34**, 21708–21719 (2021)
 451. Leonard, M.: Promotional analysis and forecasting for demand planning: a practical time series approach. with exhibits **1** (2001)
 452. Lester, B., Al-Rfou, R., Constant, N.: The power of scale for parameter-efficient prompt tuning. In: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language*

- Processing, pp. 3045–3059. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic (2021)
453. Levine, Y., Wies, N., Sharir, O., Bata, H., Shashua, A.: Limits to depth efficiencies of self-attention. *Advances in Neural Information Processing Systems* **33**, 22640–22651 (2020)
 454. Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., Zettlemoyer, L.: Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 7871–7880 (2020)
 455. Li, A.C., Prabhudesai, M., Duggal, S., Brown, E., Pathak, D.: Your diffusion model is secretly a zero-shot classifier. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 2206–2217 (2023)
 456. Li, B., Qi, X., Lukasiewicz, T., Torr, P.: Controllable text-to-image generation. *NeurIPS* (2019)
 457. Li, C., Zhuang, B., Wang, G., Liang, X., Chang, X., Yang, Y.: Automated progressive learning for efficient training of vision transformers. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12486–12496 (2022)
 458. Li, H., Guo, D., Fan, W., Xu, M., Song, Y.: Multi-step jailbreaking privacy attacks on chatgpt. *arXiv preprint arXiv:2304.05197* (2023)
 459. Li, H., Moon, J.T., Purkayastha, S., Celi, L.A., Trivedi, H., Gichoya, J.W.: Ethics of large language models in medicine and medical research. *The Lancet Digital Health* **5**(6), e333–e335 (2023)
 460. Li, H., Moon, J.T., Purkayastha, S., Celi, L.A., Trivedi, H., Gichoya, J.W.: Ethics of large language models in medicine and medical research. *The Lancet Digital Health* **5**(6), e333–e335 (2023)
 461. Li, H., Wang, M., Liu, S., Chen, P.Y.: A theoretical understanding of shallow vision transformers: Learning, generalization, and sample complexity. In: *International Conference on Learning Representations* (2023)
 462. Li, H., Wang, M., Liu, S., Chen, P.Y., Xiong, J.: Generalization guarantee of training graph convolutional networks with graph topology sampling. In: *International Conference on Machine Learning*, pp. 13014–13051. PMLR (2022)
 463. Li, H., Wang, M., Lu, S., Cui, X., Chen, P.Y.: How do nonlinear transformers acquire generalization-guaranteed cot ability? In: *High-dimensional Learning Dynamics 2024: The Emergence of Structure and Reasoning* (2024). URL <https://openreview.net/forum?id=8pM8IrT6Xo>
 464. Li, H., Wang, M., Lu, S., Cui, X., Chen, P.Y.: How do nonlinear transformers learn and generalize in in-context learning? In: *International Conference on Machine Learning* (2024)
 465. Li, H., Wang, M., Ma, T., Liu, S., ZHANG, Z., Chen, P.Y.: What improves the generalization of graph transformer? a theoretical dive into self-attention and positional encoding. In: *NeurIPS 2023 Workshop: New Frontiers in Graph Learning* (2023). URL <https://openreview.net/forum?id=BaxFC3z9R6>
 466. Li, H., Wang, M., Ma, T., Liu, S., Zhang, Z., Chen, P.Y.: What improves the generalization of graph transformers? a theoretical dive into the self-attention and positional encoding. *International Conference on Machine Learning* (2024)
 467. Li, H., Wang, M., Zhang, S., Liu, S., Chen, P.Y.: Learning on transformers is provable low-rank and sparse: A one-layer analysis. *arXiv preprint arXiv:2406.17167* (2024)
 468. Li, H., Xu, Z., Taylor, G., Studer, C., Goldstein, T.: Visualizing the loss landscape of neural nets. In: *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3–8, 2018, Montréal, Canada*, pp. 6391–6401 (2018)
 469. Li, H., Zhang, S., Wang, M.: Learning and generalization of one-hidden-layer neural networks, going beyond standard gaussian data. In: *2022 56th Annual Conference on Information Sciences and Systems (CISS)*, pp. 37–42. IEEE (2022)
 470. Li, H., Zhang, Y., Koto, F., Yang, Y., Zhao, H., Gong, Y., Duan, N., Baldwin, T.: Cmmlu: Measuring massive multitask language understanding in chinese (2023)

471. Li, J., Cheng, X., Zhao, W.X., Nie, J.Y., Wen, J.R.: Halueval: A large-scale hallucination evaluation benchmark for large language models. arXiv e-prints pp. arXiv-2305 (2023)
472. Li, J., Deng, L., Gong, Y., Haeb-Umbach, R.: An overview of noise-robust automatic speech recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* **22**(4), 745–777 (2014)
473. Li, L., Jamieson, K., Rostamizadeh, A., Gonina, E., Hardt, M., Recht, B., Talwalkar, A.: Massively parallel hyperparameter tuning. arXiv preprint arXiv:1810.05934 **5** (2018)
474. Li, L.H., Zhang, P., Zhang, H., Yang, J., Li, C., Zhong, Y., Wang, L., Yuan, L., Zhang, L., Hwang, J.N., et al.: Grounded language-image pre-training. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10965–10975 (2022)
475. Li, M., Ruan, W., Liu, X., Soldaini, L., Hamza, W., Su, C.: Improving spoken language understanding by exploiting asr n-best hypotheses. arXiv preprint arXiv:2001.05284 (2020)
476. Li, M., Zhao, Y., Yu, B., Song, F., Li, H., Yu, H., Li, Z., Huang, F., Li, Y.: Api-bank: A comprehensive benchmark for tool-augmented llms (2023)
477. Li, N., Arnold, D.M., Down, D.G., Barty, R., Blake, J., Chiang, F., Courtney, T., Waito, M., Trifunov, R., Heddle, N.M.: From demand forecasting to inventory ordering decisions for red blood cells through integrating machine learning, statistical modeling, and inventory optimization. *Transfusion* **62**(1), 87–99 (2022)
478. Li, N., Pan, A., Gopal, A., Yue, S., Berrios, D., Gatti, A., Li, J.D., Dombrowski, A.K., Goel, S., Phan, L., et al.: The wmdp benchmark: Measuring and reducing malicious use with unlearning. arXiv preprint arXiv:2403.03218 (2024)
479. Li, X., Feng, J., Meng, Y., Han, Q., Wu, F., Li, J.: A unified mrc framework for named entity recognition. In: *Annual Meeting of the Association for Computational Linguistics (ACL)* (2020)
480. Li, X., Tramer, F., Liang, P., Hashimoto, T.: Large language models can be strong differentially private learners. In: *International Conference on Learning Representations* (2021)
481. Li, X., Zhang, T., Dubois, Y., Taori, R., Gulrajani, I., Guestrin, C., Liang, P., Hashimoto, T.B.: AlpacaEval: An automatic evaluator of instruction-following models. https://github.com/tatsu-lab/alpaca_eval (2023)
482. Li, X., Zhu, X., Ma, Z., Liu, X., Shah, S.: Are chatgpt and gpt-4 general-purpose solvers for financial text analytics? an examination on several typical tasks. arXiv preprint arXiv:2305.05862 (2023)
483. Li, X.L., Liang, P.: Prefix-tuning: Optimizing continuous prompts for generation. arXiv preprint arXiv:2101.00190 (2021)
484. Li, X.L., Liang, P.: Prefix-tuning: Optimizing continuous prompts for generation (2021)
485. Li, X.L., Thackstun, J., Gulrajani, I., Liang, P., Hashimoto, T.B.: Diffusion-lm improves controllable text generation. In: *ArXiv* (2022)
486. Li, Y., Ildiz, M.E., Papailiopoulos, D., Oymak, S.: Transformers as algorithms: Generalization and stability in in-context learning. In: *International Conference on Machine Learning* (2023)
487. Li, Y., Li, Y., Risteski, A.: How do transformers learn topic structure: Towards a mechanistic understanding. arXiv preprint arXiv:2303.04245 (2023)
488. Li, Y., Li, Z., Zhang, K., Dan, R., Jiang, S., Zhang, Y.: Chatdoctor: A medical chat model fine-tuned on a large language model meta-ai (llama) using medical domain knowledge. *Cureus* **15**(6) (2023)
489. Li, Y., Liang, Y.: Learning overparameterized neural networks via stochastic gradient descent on structured data. In: *Advances in Neural Information Processing Systems*, pp. 8157–8166 (2018)
490. Li, Y., Tsai, Y.L., Yu, C.M., Chen, P.Y., Ren, X.: Exploring the benefits of visual prompting in differential privacy. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 5158–5167 (2023)

491. Li, Z., Yang, T., Wang, P., Cheng, J.: Q-vit: Fully differentiable quantization for vision transformer. arXiv preprint arXiv:2201.07703 (2022)
492. Liang, P., Bommasani, R., Lee, T., Tsipras, D., Soylu, D., Yasunaga, M., Zhang, Y., Narayanan, D., Wu, Y., Kumar, A., et al.: Holistic evaluation of language models. arXiv preprint arXiv:2211.09110 (2022)
493. Liang, W., Yüsekçönül, M., Mao, Y., Wu, E., Zou, J.: GPT detectors are biased against non-native english writers. CoRR **abs/2304.02819** (2023)
494. Liang, X., Song, S., Niu, S., Li, Z., Xiong, F., Tang, B., Wy, Z., He, D., Cheng, P., Wang, Z., Deng, H.: Uhgeval: Benchmarking the hallucination of chinese large language models via unconstrained generation (2023)
495. Liang, Y., Chongjian, G., Tong, Z., Song, Y., Wang, J., Xie, P.: Not all patches are what you need: Expediting vision transformers via token reorganizations. In: International Conference on Learning Representations (2022)
496. Liao, N., Shi, B., Cao, M., Zhang, X., Tian, Q., Yan, J.: Rethinking visual prompt learning as masked visual token modeling. arXiv preprint arXiv:2303.04998 (2023)
497. Liao, Q.V., Vaughan, J.W.: Ai transparency in the age of llms: A human-centered research roadmap. arXiv preprint arXiv:2306.01941 (2023)
498. Liaw, R., Liang, E., Nishihara, R., Moritz, P., Gonzalez, J.E., Stoica, I.: Tune: A research platform for distributed model selection and training. arXiv preprint arXiv:1807.05118 (2018)
499. Liesenfeld, A., Lopez, A., Dingemanse, M.: Opening up chatgpt: Tracking openness, transparency, and accountability in instruction-tuned text generators. In: Proceedings of the 5th International Conference on Conversational User Interfaces, pp. 1–6 (2023)
500. Likhoshesterov, V., Choromanski, K., Weller, A.: On the expressive power of self-attention matrices. arXiv preprint arXiv:2106.03764 (2021)
501. Lin, C.Y.: Rouge: A package for automatic evaluation of summaries. In: Text summarization branches out, pp. 74–81 (2004)
502. Lin, J., Zhao, H., Zhang, A., Wu, Y., Ping, H., Chen, Q.: Agentsims: An open-source sandbox for large language model evaluation. arXiv preprint arXiv:2308.04026 (2023)
503. Lin, S., Hilton, J., Evans, O.: Truthfulqa: Measuring how models mimic human falsehoods. arXiv preprint arXiv:2109.07958 (2021)
504. Lin, Y., Zhang, T., Sun, P., Li, Z., Zhou, S.: Fq-vit: Post-training quantization for fully quantized vision transformer (2022)
505. Lin, Z., Sun, Y., Shi, Y., Wang, X., Huang, L., Shen, L., Tao, D.: Efficient federated prompt tuning for black-box large pre-trained models. arXiv preprint arXiv:2310.03123 (2023)
506. Linardatos, P., Papastefanopoulos, V., Kotsiantis, S.: Explainable ai: A review of machine learning interpretability methods. Entropy **23**(1), 18 (2020)
507. Lindberg, S.I.: Mapping accountability: core concept and subtypes. International review of administrative sciences **79**(2), 202–226 (2013)
508. Linegar, M., Kocielnik, R., Alvarez, R.M.: Large language models and political science. Frontiers in Political Science **5**, 1257092 (2023)
509. Liu, A., Pan, L., Hu, X., Meng, S., Wen, L.: A semantic invariant robust watermark for large language models. CoRR **abs/2310.06356** (2023)
510. Liu, A., Swayamdipta, S., Smith, N.A., Choi, Y.: Wanli: Worker and ai collaboration for natural language inference dataset creation. In: Findings of the Association for Computational Linguistics: EMNLP 2022, pp. 6826–6847 (2022)
511. Liu, B., Liu, Q., Stone, P.: Continual learning and private unlearning. In: Conference on Lifelong Learning Agents, pp. 243–254. PMLR (2022)
512. Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning (2023)
513. Liu, H., Li, Z., Hall, D., Liang, P., Ma, T.: Sophia: A scalable stochastic second-order optimizer for language model pre-training. arXiv preprint arXiv:2305.14342 (2023)
514. Liu, H., Ma, Z., Yang, L., Zhou, T., Xia, R., Wang, Y., Wen, Q., Sun, L.: Sadi: A self-adaptive decomposed interpretable framework for electric load forecasting under extreme events. In: IEEE International Conference on Acoustics, Speech and Signal Processing (2023)

515. Liu, H., Ning, R., Teng, Z., Liu, J., Zhou, Q., Zhang, Y.: Evaluating the logical reasoning ability of chatgpt and gpt-4 (2023)
516. Liu, H., Sferrazza, C., Abbeel, P.: Languages are rewards: Hindsight finetuning using human feedback. arXiv preprint arXiv:2302.02676 (2023)
517. Liu, H., Tam, D., Muqeeth, M., Mohta, J., Huang, T., Bansal, M., Raffel, C.A.: Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning. *Advances in Neural Information Processing Systems* **35**, 1950–1965 (2022)
518. Liu, J., Shen, D., Zhang, Y., Dolan, W.B., Carin, L., Chen, W.: What makes good in-context examples for gpt-3? In: *Proceedings of Deep Learning Inside Out (DeeLIO 2022): The 3rd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pp. 100–114 (2022)
519. Liu, J., Xia, C.S., Wang, Y., Zhang, L.: Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation (2023)
520. Liu, L., Ren, Y., Lin, Z., Zhao, Z.: Pseudo numerical methods for diffusion models on manifolds. In: *ICLR* (2022)
521. Liu, R., Yang, R., Jia, C., Zhang, G., Zhou, D., Dai, A.M., Yang, D., Vosoughi, S.: Training socially aligned language models in simulated human society. arXiv preprint arXiv:2305.16960 (2023)
522. Liu, S., Chen, P., Kailkhura, B., Zhang, G., III, A.O.H., Varshney, P.K.: A primer on zeroth-order optimization in signal processing and machine learning: Principals, recent advances, and applications. *IEEE Signal Process. Mag.* **37**(5), 43–54 (2020)
523. Liu, S., Chen, P.Y., Chen, X., Hong, M.: signSGD via zeroth-order oracle. *International Conference on Learning Representations* (2019)
524. Liu, S., Chen, P.Y., Kailkhura, B., Zhang, G., Hero, A., Varshney, P.K.: A primer on zeroth-order optimization in signal processing and machine learning. *IEEE Signal Processing Magazine* (2020)
525. Liu, S., Kailkhura, B., Chen, P.Y., Ting, P., Chang, S., Amini, L.: Zeroth-order stochastic variance reduction for nonconvex optimization. In: *Advances in Neural Information Processing Systems*, pp. 3731–3741 (2018)
526. Liu, S., Yao, Y., Jia, J., Casper, S., Baracaldo, N., Hase, P., Xu, X., Yao, Y., Li, H., Varshney, K.R., et al.: Rethinking machine unlearning for large language models. arXiv preprint arXiv:2402.08787 (2024)
527. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: Ssd: Single shot multibox detector. In: *European conference on computer vision*, pp. 21–37. Springer (2016)
528. Liu, X., Ji, K., Fu, Y., Tam, W., Du, Z., Yang, Z., Tang, J.: P-tuning: Prompt tuning can be comparable to fine-tuning across scales and tasks. In: *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 61–68 (2022)
529. Liu, X., Xu, N., Chen, M., Xiao, C.: Autodan: Generating stealthy jailbreak prompts on aligned large language models. *CoRR* **abs/2310.04451** (2023)
530. Liu, Y., Deng, G., Li, Y., Wang, K., Zhang, T., Liu, Y., Wang, H., Zheng, Y., Liu, Y.: Prompt injection attack against llm-integrated applications (2023)
531. Liu, Y., Deng, G., Xu, Z., Li, Y., Zheng, Y., Zhang, Y., Zhao, L., Zhang, T., Liu, Y.: Jailbreaking chatgpt via prompt engineering: An empirical study. arXiv preprint arXiv:2305.13860 (2023)
532. Liu, Y., Jia, Y., Geng, R., Jia, J., Gong, N.Z.: Prompt injection attacks and defenses in llm-integrated applications. arXiv preprint arXiv:2310.12815 (2023)
533. Liu, Y., Le-Cong, T., Widyasari, R., Tantithamthavorn, C., Li, L., Le, X.B.D., Lo, D.: Refining chatgpt-generated code: Characterizing and mitigating code quality issues (2023)
534. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692 (2019)

535. Liu, Y., Wu, H., Wang, J., Long, M.: Non-stationary transformers: Exploring the stationarity in time series forecasting. *Advances in Neural Information Processing Systems* **35**, 9881–9893 (2022)
536. Liu, Y., Yao, Y., Ton, J.F., Zhang, X., Cheng, R.G.H., Klochkov, Y., Taufiq, M.F., Li, H.: Trustworthy llms: a survey and guideline for evaluating large language models' alignment. *arXiv preprint arXiv:2308.05374* (2023)
537. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10012–10022 (2021)
538. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 10012–10022 (2021)
539. Liu, Z., Luo, P., Wang, X., Tang, X.: Deep learning face attributes in the wild. In: *ICCV* (2015)
540. Liu, Z., Qiao, A., Neiswanger, W., Wang, H., Tan, B., Tao, T., Li, J., Wang, Y., Sun, S., Pangarkar, O., Fan, R., Gu, Y., Miller, V., Zhuang, Y., He, G., Li, H., Koto, F., Tang, L., Ranjan, N., Shen, Z., Ren, X., Iriondo, R., Mu, C., Hu, Z., Schulze, M., Nakov, P., Baldwin, T., Xing, E.P.: Llm360: Towards fully transparent open-source llms (2023)
541. Liu, Z., Wang, J., Dao, T., Zhou, T., Yuan, B., Song, Z., Shrivastava, A., Zhang, C., Tian, Y., Re, C., et al.: Deja vu: Contextual sparsity for efficient llms at inference time. In: *International Conference on Machine Learning*, pp. 22137–22176. PMLR (2023)
542. Liu, Z., Wang, Y., Han, K., Zhang, W., Ma, S., Gao, W.: Post-training quantization for vision transformer. *Advances in Neural Information Processing Systems* **34**, 28092–28103 (2021)
543. Liu, Z., Yu, X., Zhang, L., Wu, Z., Cao, C., Dai, H., Zhao, L., Liu, W., Shen, D., Li, Q., Liu, T., Zhu, D., Li, X.: Deid-gpt: Zero-shot medical text de-identification by gpt-4 (2023)
544. Loedeman, J., Stol, M.C., Han, T., Asano, Y.M.: Prompt generation networks for efficient adaptation of frozen vision transformers. *arXiv preprint arXiv:2210.06466* (2022)
545. Lopez-Lira, A., Tang, Y.: Can chatgpt forecast stock price movements? return predictability and large language models (2023)
546. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101* (2017)
547. Loureiro, B., Sicuro, G., Gerbelot, C., Pacco, A., Krzakala, F., Zdeborová, L.: Learning gaussian mixtures with generalized linear models: Precise asymptotics in high-dimensions. *Advances in Neural Information Processing Systems* **34**, 10144–10157 (2021)
548. Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., Zhu, J.: Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. In: *NIPS* (2022)
549. Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., Zhu, J.: Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. In: *NIPS* (2022)
550. Lucy, L., Bamman, D.: Gender and representation bias in gpt-3 generated stories. In: *Proceedings of the Third Workshop on Narrative Understanding*, pp. 48–55 (2021)
551. Lund, B.D., Wang, T., Mannuru, N.R., Nie, B., Shimray, S., Wang, Z.: Chatgpt and a new academic reality: Artificial intelligence-written research papers and the ethics of the large language models in scholarly publishing. *Journal of the Association for Information Science and Technology* **74**(5), 570–581 (2023)
552. Luo, J.H., Wu, J., Lin, W.: Thinet: A filter level pruning method for deep neural network compression. In: *Proceedings of the IEEE international conference on computer vision*, pp. 5058–5066 (2017)
553. Luo, Y., Li, H., Shi, L., Wu, X.M.: Enhancing graph transformers with hierarchical distance structural encoding (2024)
554. Luo, Y., Yang, Z., Meng, F., Li, Y., Zhou, J., Zhang, Y.: An empirical study of catastrophic forgetting in large language models during continual fine-tuning. *arXiv preprint arXiv:2308.08747* (2023)

555. Luo, Y., Zhang, J., Fan, S., Yang, K., Wu, Y., Qiao, M., Nie, Z.: Biomedgpt: Open multi-modal generative pre-trained transformer for biomedicine. arXiv preprint arXiv:2308.09442 (2023)
556. Ma, J.: Imagenette classification (2021)
557. Ma, R., Gales, M.J., Knill, K., Qian, M.: N-best t5: Robust asr error correction using multiple input hypotheses and constrained decoding space. arXiv preprint arXiv:2303.00456 (2023)
558. Ma, X., Fang, G., Wang, X.: LLM-pruner: On the structural pruning of large language models. In: Thirty-seventh Conference on Neural Information Processing Systems (2023). URL <https://openreview.net/forum?id=J8Ajf9WfXP>
559. Ma, X., Yuan, G., Shen, X., Chen, T., Chen, X., Chen, X., Liu, N., Qin, M., Liu, S., Wang, Z., et al.: Sanity checks for lottery tickets: Does your winning ticket really win the jackpot? Advances in Neural Information Processing Systems **34**, 12749–12760 (2021)
560. Ma, Z., Ethayarajh, K., Thrush, T., Jain, S., Wu, L., Jia, R., Potts, C., Williams, A., Kiela, D.: Dynaboard: An evaluation-as-a-service platform for holistic next-generation benchmarking. Advances in Neural Information Processing Systems **34**, 10351–10367 (2021)
561. Ma, Z., Wu, W., Zheng, Z., Guo, Y., Chen, Q., Zhang, S., Chen, X.: Leveraging speech ptm, text llm, and emotional tts for speech emotion recognition. arXiv preprint arXiv:2309.10294 (2023)
562. Maddison, C.J., Mnih, A., Teh, Y.W.: The concrete distribution: A continuous relaxation of discrete random variables. In: ICLR (2017)
563. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. In: International Conference on Learning Representations (2018)
564. Magazine, Q.: What does it mean to align ai with human values? (2022). URL <https://www.quantamagazine.org/what-does-it-mean-to-align-ai-with-human-values-20221213/>
565. Mahajan, D., Girshick, R., Ramanathan, V., He, K., Paluri, M., Li, Y., Bharambe, A., Van Der Maaten, L.: Exploring the limits of weakly supervised pretraining. In: Proceedings of the European conference on computer vision (ECCV), pp. 181–196 (2018)
566. Maini, P., Feng, Z., Schwarzschild, A., Lipton, Z.C., Kolter, J.Z.: Tofu: A task of fictitious unlearning for llms (2024)
567. Makridakis, S., Spiliotis, E., Assimakopoulos, V.: The m4 competition: Results, findings, conclusion and way forward. International Journal of Forecasting **34**(4), 802–808 (2018)
568. Malladi, S., Gao, T., Nichani, E., Damian, A., Lee, J.D., Chen, D., Arora, S.: Fine-tuning language models with just forward passes. arXiv preprint arXiv:2305.17333 (2023)
569. Mani, A., Palaskar, S., Meripo, N.V., Konam, S., Metze, F.: Asr error correction and domain adaptation using machine translation. In: ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 6344–6348. IEEE (2020)
570. de Masson d’Autume, C., Mohamed, S., Rosca, M., Rae, J.W.: Training language gans from scratch. In: NeurIPS, pp. 4302–4313 (2019)
571. Mattern, J., Jin, Z., Weggenmann, B., Schoelkopf, B., Sachan, M.: Differentially private language models for secure data sharing. In: Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, pp. 4860–4873 (2022)
572. Mattern, J., Weggenmann, B., Kerschbaum, F.: The limits of word level differential privacy. In: Findings of the Association for Computational Linguistics: NAACL 2022, pp. 867–881 (2022)
573. Maus, N., Chao, P., Wong, E., Gardner, J.R.: Black box adversarial prompting for foundation models. In: The Second Workshop on New Frontiers in Adversarial Machine Learning (2023)
574. McKenna, N., Li, T., Cheng, L., Hosseini, M.J., Johnson, M., Steedman, M.: Sources of hallucination by large language models on inference tasks (2023)
575. McSherry, F., Talwar, K.: Mechanism design via differential privacy. In: 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS’07), pp. 94–103. IEEE (2007)

576. Meehan, C., Mrini, K., Chaudhuri, K.: Sentence-level privacy for document embeddings. In: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 3367–3380 (2022)
577. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., Galstyan, A.: A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)* **54**(6), 1–35 (2021)
578. Mehrotra, A., Zampetakis, M., Kassianik, P., Nelson, B., Anderson, H., Singer, Y., Karbasi, A.: Tree of attacks: Jailbreaking black-box llms automatically. *CoRR* **abs/2312.02119** (2023)
579. Mei, K., Patel, V.M.: VIDM: video implicit diffusion models. *CoRR* **abs/2212.00235** (2022)
580. Merity, S., Xiong, C., Bradbury, J., Socher, R.: Pointer sentinel mixture models (2016)
581. Meta: Ai at meta (2023). <https://ai.meta.com>
582. Meta: Responsible use guide: your resource for building responsibly (2023). URL <https://ai.meta.com/llama/responsible-use-guide/>
583. Meyer, J.G., Urbanowicz, R.J., Martin, P.C., O'Connor, K., Li, R., Peng, P.C., Bright, T.J., Tatonetti, N., Won, K.J., Gonzalez-Hernandez, G., et al.: Chatgpt and large language models in academia: opportunities and challenges. *BioData Mining* **16**(1), 20 (2023)
584. Microsoft: “Introducing the new Bing.”. [tps://www.bing.com/new](https://www.bing.com/new) (2023). [Online; accessed 4-Apr-2023]
585. Mignacco, F., Krzakala, F., Lu, Y., Urbani, P., Zdeborova, L.: The role of regularization in classification of high-dimensional noisy gaussian mixture. In: International conference on machine learning, pp. 6874–6883. PMLR (2020)
586. Mikolov, T., Karafiát, M., Burget, L., Cernocký, J., Khudanpur, S.: Recurrent neural network based language model. In: Interspeech, vol. 2, pp. 1045–1048. Makuhari (2010)
587. MintMesh: Large language models: The future of b2b software (2023). URL <https://www.mintmesh.ai/blog/large-language-models-the-future-of-b2b-software#:~:text=From%20refining%20customer%20support%20to,era%20of%20efficiency%20and%20innovation>
588. Mirchandani, S., Xia, F., Florence, P., Driess, D., Arenas, M.G., Rao, K., Sadigh, D., Zeng, A., et al.: Large language models as general pattern machines. In: Proceedings of the 7th Annual Conference on Robot Learning (2023)
589. Mireshghallah, F., Inan, H.A., Hasegawa, M., Rühle, V., Berg-Kirkpatrick, T., Sim, R.: Privacy regularization: Joint privacy-utility optimization in language models. *arXiv preprint arXiv:2103.07567* (2021)
590. Misra, D., Goyal, A., Runwal, B., Chen, P.Y.: Reprogramming under constraints: Revisiting efficient and reliable transferability of lottery tickets. *arXiv preprint arXiv:2308.14969* (2023)
591. Mitchell, E., Lee, Y., Khazatsky, A., Manning, C.D., Finn, C.: Detectgpt: Zero-shot machine-generated text detection using probability curvature. *CoRR* **abs/2301.11305** (2023)
592. Müller, V.C.: Ethics of Artificial Intelligence and Robotics. In: E.N. Zalta, U. Nodelman (eds.) *The Stanford Encyclopedia of Philosophy*, Fall 2023 edn. Metaphysics Research Lab, Stanford University (2023)
593. Mo, L., Wang, B., Chen, M., Sun, H.: How trustworthy are open-source llms? an assessment under malicious demonstrations shows their vulnerabilities (2023)
594. Molchanov, P., Tyree, S., Karras, T., Aila, T., Kautz, J.: Pruning convolutional neural networks for resource efficient inference. In: International Conference on Learning Representations (2016)
595. Montagna, S., Ferretti, S., Klopfenstein, L.C., Florio, A., Pengo, M.F.: Data decentralisation of llm-based chatbot systems in chronic disease self-management. In: Proceedings of the 2023 ACM Conference on Information Technology for Social Good, pp. 205–212 (2023)
596. Moor, J.H.: The nature, importance, and difficulty of machine ethics. *IEEE intelligent systems* **21**(4), 18–21 (2006)
597. Moosavi-Dezfooli, S.M., Fawzi, A., Frossard, P.: Deepfool: a simple and accurate method to fool deep neural networks. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 2574–2582 (2016)

598. Mosbach, M., Pimentel, T., Ravfogel, S., Klakow, D., Elazar, Y.: Few-shot fine-tuning vs. in-context learning: A fair comparison and evaluation. *arXiv preprint arXiv:2305.16938* (2023)
599. Motoki, F., Neto, V.P., Rodrigues, V.: More human than human: Measuring chatgpt political bias. *Public Choice* pp. 1–21 (2023)
600. Mowshowitz, Z.: Jailbreaking chatgpt on release day. <https://www.lesswrong.com/posts/RYcoJdvmoBbi5Nax7/jailbreaking-chatgpt-on-release-day> (2022)
601. Mulgan, R.: ‘accountability’: an ever-expanding concept? *Public administration* **78**(3), 555–573 (2000)
602. Nakano, R., Hilton, J., Balaji, S., Wu, J., Ouyang, L., Kim, C., Hesse, C., Jain, S., Kosaraju, V., Saunders, W., et al.: Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332* (2021)
603. Nascimento, C., Pimentel, A.: Do large language models understand chemistry? a conversation with. *Journal of Chemical Information and Modeling* **63** (2023). <https://doi.org/10.1021/acs.jcim.3c00285>
604. Nay, J.J., Karamardian, D., Lawsky, S.B., Tao, W., Bhat, M., Jain, R., Lee, A.T., Choi, J.H., Kasai, J.: Large language models as tax attorneys: A case study in legal capabilities emergence (2023)
605. Near, J.: Differential privacy at scale: Uber and Berkeley collaboration. In: *Enigma 2018* (Enigma 2018) (2018)
606. Nesterov, Y., Spokoiny, V.: Random gradient-free minimization of convex functions. *Foundations of Computational Mathematics* **17**(2), 527–566 (2017)
607. Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., Ng, A.Y.: Reading digits in natural images with unsupervised feature learning. *NIPS Workshop on Deep Learning and Unsupervised Feature Learning* (2011)
608. Neyshabur, B., Sedghi, H., Zhang, C.: What is being transferred in transfer learning? *Advances in neural information processing systems* **33**, 512–523 (2020)
609. Nguyen, T.T., Huynh, T.T., Nguyen, P.L., Liew, A.W.C., Yin, H., Nguyen, Q.V.H.: A survey of machine unlearning. *arXiv preprint arXiv:2209.02299* (2022)
610. Nichol, A.Q., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: GLIDE: towards photorealistic image generation and editing with text-guided diffusion models. In: *ICML* (2022)
611. Nie, Y., Nguyen, N.H., Sinthong, P., Kalagnanam, J.: A time series is worth 64 words: Long-term forecasting with transformers. In: *International Conference on Learning Representations* (2023)
612. Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes. In: *Proceedings of the Indian Conference on Computer Vision, Graphics and Image Processing* (2008)
613. Nissenbaum, H.: Accountability in a computerized society. *Science and engineering ethics* **2**, 25–42 (1996)
614. Novelli, C., Taddeo, M., Floridi, L.: Accountability in artificial intelligence: what it is and how it works. *AI & SOCIETY* pp. 1–12 (2023)
615. Oh, C., Hwang, H., Lee, H.y., Lim, Y., Jung, G., Jung, J., Choi, H., Song, K.: Blackvip: Black-box visual prompting for robust transfer learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 24224–24235 (2023)
616. Okolo, G.I., Katsigiannis, S., Ramzan, N.: Ievit: An enhanced vision transformer architecture for chest x-ray image classification. *Computer Methods and Programs in Biomedicine* **226**, 107141 (2022)
617. OpenAI: Chatgpt: Optimizing language models for dialogue (2022). URL <https://openai.com/blog/chatgpt>
618. OpenAI: ChatGPT plugins. <https://openai.com/blog/chatgpt-plugins> (2023). [Online; accessed 16-Apr-2023]
619. OpenAI: GPT-4 technical report. *CoRR* **abs/2303.08774** (2023)
620. OpenAI: Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023)
621. OpenAI: Gpt-4 technical report (2023)

622. OpenAI: GPT-4V(ision) system card. <https://openai.com/research/gpt-4v-system-card> (2023)
623. OpenAI: Openai (2023). <https://www.openai.com>
624. Oreshkin, B.N., Carpo, D., Chapados, N., Bengio, Y.: N-beats: Neural basis expansion analysis for interpretable time series forecasting. In: International Conference on Learning Representations (2020)
625. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al.: Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems* **35**, 27730–27744 (2022)
626. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al.: Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems* **35**, 27730–27744 (2022)
627. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C.L., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P., Leike, J., Lowe, R.: Training language models to follow instructions with human feedback (2022)
628. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C.L., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P.F., Leike, J., Lowe, R.: Training language models to follow instructions with human feedback. In: *NeurIPS* (2022)
629. Ouyang, S., Zhang, Z., Yan, B., Liu, X., Han, J., Qin, L.: Structured chemistry reasoning with large language models. *arXiv preprint arXiv:2311.09656* (2023)
630. Oymak, S., Rawat, A.S., Soltanolkotabi, M., Thrampoulidis, C.: On the role of attention in prompt-tuning. *arXiv preprint arXiv:2306.03435* (2023)
631. Pal, S., Bhattacharya, M., Lee, S.S., Chakraborty, C.: A domain-specific next-generation large language model (llm) or chatgpt is required for biomedical engineering and research. *Annals of Biomedical Engineering* pp. 1–4 (2023)
632. Pallagani, V., Muppasani, B., Murugesan, K., Rossi, F., Srivastava, B., Horesh, L., Fabiano, F., Loreggia, A.: Understanding the capabilities of large language models for automated planning (2023)
633. Pan, A., Bhatia, K., Steinhardt, J.: The effects of reward misspecification: Mapping and mitigating misaligned models. *arXiv preprint arXiv:2201.03544* (2022)
634. Pan, L., Albalak, A., Wang, X., Wang, W.Y.: Logic-llm: Empowering large language models with symbolic solvers for faithful logical reasoning (2023)
635. Pan, Y., Pan, L., Chen, W., Nakov, P., Kan, M.Y., Wang, W.Y.: On the risk of misinformation pollution with large language models. *arXiv preprint arXiv:2305.13661* (2023)
636. Pan, Z., Zhuang, B., Liu, J., He, H., Cai, J.: Scalable vision transformers with hierarchical pooling. In: *Proceedings of the IEEE/cvf international conference on computer vision*, pp. 377–386 (2021)
637. Panayotov, V., Chen, G., Povey, D., Khudanpur, S.: Librispeech: an asr corpus based on public domain audio books. In: 2015 IEEE international conference on acoustics, speech and signal processing (ICASSP), pp. 5206–5210. IEEE (2015)
638. Pandey, A., Liu, C., Wang, Y., Saraf, Y.: Dual application of speech enhancement for automatic speech recognition. In: 2021 IEEE Spoken Language Technology Workshop (SLT), pp. 223–228. IEEE (2021)
639. Papernot, N., Abadi, M., Úlfar Erlingsson, Goodfellow, I., Talwar, K.: Semi-supervised knowledge transfer for deep learning from private training data. In: International Conference on Learning Representations (ICLR) (2017)
640. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: Bleu: a method for automatic evaluation of machine translation. In: *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pp. 311–318 (2002)
641. Parkhi, O.M., Vedaldi, A., Zisserman, A., Jawahar, C.: Cats and dogs. In: 2012 IEEE conference on computer vision and pattern recognition, pp. 3498–3505. IEEE (2012)

642. Pass, G., Chowdhury, A., Torgeson, C.: A picture of search. In: Proceedings of the 1st international conference on Scalable information systems, pp. 1–es (2006)
643. Paul, S., Chen, P.Y.: Vision transformers are robust learners. Proceedings of the AAAI Conference on Artificial Intelligence (2022)
644. Paul, S., Chen, P.Y.: Vision transformers are robust learners. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 36, pp. 2071–2081 (2022)
645. Pearce, T., Rashid, T., Kanervisto, A., Bignell, D., Sun, M., Georgescu, R., Macua, S.V., Tan, S.Z., Momennejad, I., Hofmann, K., Devlin, S.: Imitating human behaviour with diffusion models. In: CoRR (2023)
646. Penedo, G., Malartic, Q., Hesslow, D., Cojocaru, R., Cappelli, A., Alobeidli, H., Pannier, B., Almazrouei, E., Launay, J.: The refinedweb dataset for falcon llm: outperforming curated corpora with web data, and web data only. arXiv preprint arXiv:2306.01116 (2023)
647. Peng, A., Wu, M., Allard, J., Kilpatrick, L., Heidel, S.: Gpt-3.5 turbo fine-tuning and api updates (2023). URL <https://openai.com/blog/gpt-3-5-turbo-fine-tuning-and-api-updates>
648. Peng, F., Roy, S., Shahshahani, B., Beaufays, F.: Search results based n-best hypothesis rescoring with maximum entropy classification. In: 2013 IEEE Workshop on Automatic Speech Recognition and Understanding, pp. 422–427. IEEE (2013)
649. Perez, E., Huang, S., Song, F., Cai, T., Ring, R., Aslanides, J., Glaese, A., McAleese, N., Irving, G.: Red teaming language models with language models. arXiv preprint arXiv:2202.03286 (2022)
650. Perez, E., Ringer, S., Lukošiušė, K., Nguyen, K., Chen, E., Heiner, S., Pettit, C., Olsson, C., Kundu, S., Kadavath, S., et al.: Discovering language model behaviors with model-written evaluations. arXiv preprint arXiv:2212.09251 (2022)
651. Petridis, S., Perantonis, S.J.: On the relation between discriminant analysis and mutual information for supervised linear feature extraction. Pattern Recognition **37**(5), 857–874 (2004)
652. Petroni, F., Rocktäschel, T., Lewis, P., Bakhtin, A., Wu, Y., Miller, A.H., Riedel, S.: Language models as knowledge bases? arXiv preprint arXiv:1909.01066 (2019)
653. Phuong, M., Hutter, M.: Formal algorithms for transformers. arXiv preprint arXiv:2207.09238 (2022)
654. Phute, M., Helbling, A., Hull, M., Peng, S., Szyller, S., Cornelius, C., Chau, D.H.: Llm self defense: By self examination, llms know they are being tricked (2023)
655. Piet, J., Sitawarin, C., Fang, V., Mu, N., Wagner, D.: Mark my words: Analyzing and evaluating language model watermarks. arXiv preprint arXiv:2312.00273 (2023)
656. Pillutla, K., Swayamdipta, S., Zellers, R., Thickstun, J., Welleck, S., Choi, Y., Harchaoui, Z.: Mauve: Measuring the gap between neural text and human text using divergence frontiers. In: M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, J.W. Vaughan (eds.) Advances in Neural Information Processing Systems, vol. 34, pp. 4816–4828. Curran Associates, Inc. (2021). URL https://proceedings.neurips.cc/paper_files/paper/2021/file/260c2432a0eccc28ce03c10dadcd078a4-Paper.pdf
657. Pinkney, J.N.M.: Pokemon blip captions. <https://huggingface.co/datasets/lambdalabs/pokemon-blip-captions/> (2022)
658. Poggio, T., Girosi, F.: Networks for approximation and learning. Proceedings of the IEEE **78**(9), 1481–1497 (1990)
659. Popov, V., Vovk, I., Gogoryan, V., Sadekova, T., Kudinov, M.A.: Grad-tts: A diffusion probabilistic model for text-to-speech. In: ICML (2021)
660. Porsdam Mann, S., Earp, B.D., Möller, N., Vynn, S., Savulescu, J.: Autogen: A personalized large language model for academic enhancement—ethics and proof of principle. The American Journal of Bioethics pp. 1–14 (2023)
661. Prasad, A., Jyothi, P., Velmurugan, R.: An investigation of end-to-end models for robust speech recognition. In: ICASSP 2021–2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 6893–6897. IEEE (2021)
662. Preoțiuc-Pietro, D., Lamos, V., Aletras, N.: An analysis of the user occupational class through twitter content. In: Proceedings of the 53rd Annual Meeting of the Association for

- Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pp. 1754–1764 (2015)
663. Pruksachatkun, Y., Phang, J., Liu, H., Htut, P.M., Zhang, X., Pang, R.Y., Vania, C., Kann, K., Bowman, S.: Intermediate-task transfer learning with pretrained language models: When and why does it work? In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 5231–5247 (2020)
 664. Pu, X., Gao, M., Wan, X.: Summarization is (almost) dead (2023)
 665. Qi, X., Huang, K., Panda, A., Henderson, P., Wang, M., Mittal, P.: Visual adversarial examples jailbreak aligned large language models (2023)
 666. Qi, X., Zeng, Y., Xie, T., Chen, P.Y., Jia, R., Mittal, P., Henderson, P.: Fine-tuning aligned language models compromises safety, even when users do not intend to! arXiv preprint arXiv:2310.03693 (2023)
 667. Qiao, F., Zhao, L., Peng, X.: Learning to learn single domain generalization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 12556–12565 (2020)
 668. Qin, C., Zhang, A., Zhang, Z., Chen, J., Yasunaga, M., Yang, D.: Is chatgpt a general-purpose natural language processing task solver? (2023)
 669. Qin, Y., Hu, S., Lin, Y., Chen, W., Ding, N., Cui, G., Zeng, Z., Huang, Y., Xiao, C., Han, C., Fung, Y.R., Su, Y., Wang, H., Qian, C., Tian, R., Zhu, K., Liang, S., Shen, X., Xu, B., Zhang, Z., Ye, Y., Li, B., Tang, Z., Yi, J., Zhu, Y., Dai, Z., Yan, L., Cong, X., Lu, Y., Zhao, W., Huang, Y., Yan, J., Han, X., Sun, X., Li, D., Phang, J., Yang, C., Wu, T., Ji, H., Liu, Z., Sun, M.: Tool learning with foundation models (2023)
 670. Qin, Y., Liang, S., Ye, Y., Zhu, K., Yan, L., Lu, Y., Lin, Y., Cong, X., Tang, X., Qian, B., Zhao, S., Hong, L., Tian, R., Xie, R., Zhou, J., Gerstein, M., Li, D., Liu, Z., Sun, M.: Toolllm: Facilitating large language models to master 16000+ real-world apis (2023)
 671. Qiu, H., Zhang, S., Li, A., He, H., Lan, Z.: Latent jailbreak: A benchmark for evaluating text safety and output robustness of large language models. arXiv preprint arXiv:2307.08487 (2023)
 672. Qiu, H., Zhang, S., Li, A., He, H., Lan, Z.: Latent jailbreak: A test suite for evaluating both text safety and output robustness of large language models (2023)
 673. Quiring, E., Maier, A., Rieck, K.: Misleading authorship attribution of source code using adversarial learning. In: 28th USENIX Security Symposium (USENIX Security 19), pp. 479–496 (2019)
 674. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning, pp. 8748–8763. PMLR (2021)
 675. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning, pp. 8748–8763. PMLR (2021)
 676. Radford, A., Kim, J.W., Xu, T., Brockman, G., McLeavey, C., Sutskever, I.: Robust speech recognition via large-scale weak supervision. arXiv preprint arXiv:2212.04356 (2022)
 677. Radford, A., Kim, J.W., Xu, T., Brockman, G., McLeavey, C., Sutskever, I.: Robust speech recognition via large-scale weak supervision. In: International Conference on Machine Learning, pp. 28492–28518. PMLR (2023)
 678. Radford, A., Narasimhan, K., Salimans, T., Sutskever, I., et al.: Improving language understanding by generative pre-training (2018)
 679. Radford, A., Narasimhan, K., Salimans, T., Sutskever, I., et al.: Improving language understanding by generative pre-training (2018)
 680. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. OpenAI blog **1**(8), 9 (2019)
 681. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. OpenAI blog **1**(8), 9 (2019)

682. Rafailov, R., Sharma, A., Mitchell, E., Manning, C.D., Ermon, S., Finn, C.: Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems* **36** (2023)
683. Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., Liu, P.J.: Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR* (2020)
684. Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., Liu, P.J.: Exploring the limits of transfer learning with a unified text-to-text transformer (2023)
685. Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., Liu, P.J., et al.: Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.* (2020)
686. Rajapakse, J.C., Giedd, J.N., Rapoport, J.L.: Statistical approach to segmentation of single-channel cerebral mr images. *IEEE transactions on medical imaging* **16**(2), 176–186 (1997)
687. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. In: *ArXiv* (2022)
688. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125* **1**(2), 3 (2022)
689. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: *International Conference on Machine Learning*, pp. 8821–8831. *PMLR* (2021)
690. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: *International Conference on Machine Learning*, pp. 8821–8831. *PMLR* (2021)
691. Rando, J., Paleka, D., Lindner, D., Heim, L., Tramer, F.: Red-teaming the stable diffusion safety filter. In: *NeurIPS ML Safety Workshop* (2022)
692. Rando, J., Paleka, D., Lindner, D., Heim, L., Tramer, F.: Red-teaming the stable diffusion safety filter. In: *NeurIPS ML Safety Workshop* (2022)
693. Rao, A., Vashistha, S., Naik, A., Aditya, S., Choudhury, M.: Tricking llms into disobedience: Understanding, analyzing, and preventing jailbreaks. *arXiv preprint arXiv:2305.14965* (2023)
694. Rao, J.R., Rohatgi, P., et al.: Can pseudonymity really guarantee privacy? In: *USENIX Security Symposium*, pp. 85–96 (2000)
695. Rao, Y., Zhao, W., Liu, B., Lu, J., Zhou, J., Hsieh, C.J.: Dynamicvit: Efficient vision transformers with dynamic token sparsification. *Advances in neural information processing systems* **34**, 13937–13949 (2021)
696. Redmon, J., Farhadi, A.: Yolo9000: better, faster, stronger. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7263–7271 (2017)
697. Reed, S., Zolna, K., Parisotto, E., Colmenarejo, S.G., Novikov, A., Barth-Maron, G., Gimenez, M., Sulsky, Y., Kay, J., Springenberg, J.T., et al.: A generalist agent. *arXiv preprint arXiv:2205.06175* (2022)
698. Refinetti, M., Goldt, S., Krzakala, F., Zdeborová, L.: Classifying high-dimensional gaussian mixtures: Where kernel methods fail and neural networks succeed. In: *International Conference on Machine Learning*, pp. 8936–8947. *PMLR* (2021)
699. Reimers, N., Gurevych, I.: Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084* (2019)
700. Ren, J., Xu, H., Liu, Y., Cui, Y., Wang, S., Yin, D., Tang, J.: A robust semantics-based watermark for large language model against paraphrasing. *CoRR* **abs/2311.08721** (2023). <https://doi.org/10.48550/arXiv.2311.08721>
701. Ren, M., Kornblith, S., Liao, R., Hinton, G.: Scaling forward gradient with local losses. *arXiv preprint arXiv:2210.03310* (2022)
702. Rezende, D.J., Mohamed, S.: Variational inference with normalizing flows. In: *ICML* (2015)
703. Riba, E., Mishkin, D., Ponsa, D., Rublee, E., Bradski, G.: Kornia: an open source differentiable computer vision library for pytorch. In: *Winter Conference on Applications of Computer Vision* (2020)

704. Robey, A., Wong, E., Hassani, H., Pappas, G.J.: Smoothllm: Defending large language models against jailbreaking attacks. *CoRR* **abs/2310.03684** (2023)
705. Rodriguez, J., Hay, T., Gros, D., Shamsi, Z., Srinivasan, R.: Cross-domain detection of gpt-2-generated technical text. In: *NAACL*, pp. 1213–1233 (2022)
706. Roemmele, M., Bejan, C.A., Gordon, A.S.: Choice of plausible alternatives: An evaluation of commonsense causal reasoning. In: *2011 AAAI Spring Symposium Series* (2011)
707. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *CVPR* (2021)
708. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: Compvis/stable diffusion v1-4. <https://huggingface.co/CompVis/stable-diffusion-v1-4> (2022)
709. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *CVPR* (2022)
710. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *CVPR* (2022)
711. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. *arxiv*. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2022)
712. Röttger, P., Kirk, H.R., Vidgen, B., Attanasio, G., Bianchi, F., Hovy, D.: Xstest: A test suite for identifying exaggerated safety behaviours in large language models. *arXiv preprint arXiv:2308.01263* (2023)
713. Rozière, B., Gehring, J., Gloeckle, F., Sootla, S., Gat, I., Tan, X.E., Adi, Y., Liu, J., Remez, T., Rapin, J., Kozhevnikov, A., Evtimov, I., Bitton, J., Bhatt, M., Canton-Ferrer, C., Grattafiori, A., Xiong, W., Défossez, A., Copet, J., Azhar, F., Touvron, H., Martin, L., Usunier, N., Scialom, T., Synnaeve, G.: Code llama: Open foundation models for code. *CoRR* **abs/2308.12950** (2023). <https://doi.org/10.48550/arXiv.2308.12950>
714. Roziere, B., Gehring, J., Gloeckle, F., Sootla, S., Gat, I., Tan, X.E., Adi, Y., Liu, J., Remez, T., Rapin, J., et al.: Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950* (2023)
715. Rozière, B., Gehring, J., Gloeckle, F., Sootla, S., Gat, I., Tan, X.E., Adi, Y., Liu, J., Remez, T., Rapin, J., Kozhevnikov, A., Evtimov, I., Bitton, J., Bhatt, M., Ferrer, C.C., Grattafiori, A., Xiong, W., Défossez, A., Copet, J., Azhar, F., Touvron, H., Martin, L., Usunier, N., Scialom, T., Synnaeve, G.: Code llama: Open foundation models for code (2023). URL <https://ai.meta.com/research/publications/code-llama-open-foundation-models-for-code/>
716. Ruan, Y., Dubois, Y., Maddison, C.J.: Optimal representations for covariate shift. In: *International Conference on Learning Representations* (2021)
717. Rubin, O., Herzig, J., Berant, J.: Learning to retrieve prompts for in-context learning. In: *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 2655–2671 (2022)
718. Sadasivan, V.S., Kumar, A., Balasubramanian, S., Wang, W., Feizi, S.: Can ai-generated text be reliably detected? *CoRR* **abs/2303.11156** (2023)
719. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E., Ghasemipour, S.K.S., Ayan, B.K., Mahdavi, S.S., Lopes, R.G., Salimans, T., Ho, J., Fleet, D.J., Norouzi, M.: Photorealistic text-to-image diffusion models with deep language understanding. In: *ArXiv* (2022)
720. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems (NeurIPS)* **35**, 36479–36494 (2022)
721. Sakaguchi, K., Bras, R.L., Bhagavatula, C., Choi, Y.: Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM* **64**(9), 99–106 (2021)
722. Salazar, J., Liang, D., Nguyen, T.Q., Kirchhoff, K.: Masked language model scoring. *arXiv preprint arXiv:1910.14659* (2019)
723. Salimans, T., Goodfellow, I.J., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training gans. In: *NIPS* (2016)

724. Salimans, T., Ho, J.: Progressive distillation for fast sampling of diffusion models. In: ICLR (2022)
725. Samaan, J., Yeo, Y., Rajeev, N., Hawley, L., Abel, S., Ng, W.H., Srinivasan, N., Park, J., Burch, M., Watson, R., Liran, O., Samakar, K.: Assessing the accuracy of responses by the language model chatgpt to questions regarding bariatric surgery. *Obesity Surgery* **33**, 1–7 (2023). <https://doi.org/10.1007/s11695-023-06603-5>
726. Sanh, V., Webson, A., Raffel, C., Bach, S.H., Sutawika, L., Alyafeai, Z., Chaffin, A., Stiegler, A., Scao, T.L., Raja, A., Dey, M., Bari, M.S., Xu, C., Thakker, U., Sharma, S.S., Szczechla, E., Kim, T., Chhablani, G., Nayak, N., Datta, D., Chang, J., Jiang, M.T.J., Wang, H., Manica, M., Shen, S., Yong, Z.X., Pandey, H., Bawden, R., Wang, T., Neeraj, T., Rozen, J., Sharma, A., Santilli, A., Fevry, T., Fries, J.A., Teehan, R., Bers, T., Biderman, S., Gao, L., Wolf, T., Rush, A.M.: Multitask prompted training enables zero-shot task generalization (2022)
727. Sanjay-Gopal, S., Hebert, T.J.: Bayesian pixel classification using spatially variant finite mixtures and the generalized em algorithm. *IEEE Transactions on Image Processing* **7**(7), 1014–1028 (1998)
728. Schaeffer, R., Miranda, B., Koyejo, S.: Are emergent abilities of large language models a mirage? *Advances in Neural Information Processing Systems* **36** (2024)
729. Schick, T., Schütze, H.: Exploiting cloze-questions for few-shot text classification and natural language inference. In: EACL, pp. 255–269 (2021)
730. Schick, T., Schütze, H.: Few-shot text generation with pattern-exploiting training. In: *Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP)* (2021)
731. Schneider, S.H., Dickinson, R.E.: Climate modeling. *Reviews of Geophysics* **12**(3), 447–493 (1974)
732. Schramowski, P., Brack, M., Deiseroth, B., Kersting, K.: Safe latent diffusion: Mitigating inappropriate degeneration in diffusion models. In: CVPR (2023)
733. Schramowski, P., Brack, M., Deiseroth, B., Kersting, K.: Safe latent diffusion: Mitigating inappropriate degeneration in diffusion models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 22522–22531 (2023)
734. Schramowski, P., Tauchmann, C., Kersting, K.: Can machines help us answering question 16 in datasheets, and in turn reflecting on inappropriate content? In: *ACM Conference on Fairness, Accountability, and Transparency (FAccT)* (2022)
735. Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., Schramowski, P., Kundurthy, S., Crowson, K., Schmidt, L., Kaczmarczyk, R., Jitsev, J.: LAION-5B: an open large-scale dataset for training next generation image-text models. In: NIPS (2022)
736. Schuhmann, C., Vencu, R., Beaumont, R., Kaczmarczyk, R., Mullis, C., Katta, A., Coombes, T., Jitsev, J., Komatsuzaki, A.: LAION-400M: open dataset of clip-filtered 400 million image-text pairs. NIPS Workshop (2021)
737. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347 (2017)
738. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal policy optimization algorithms. *CoRR* **abs/1707.06347** (2017)
739. Science, T.D.: Palm: Efficiently training massive language models (2023). URL <https://towardsdatascience.com/palm-efficiently-training-massive-language-models-b82d6cc1582>
740. Sefara, T.J., Mbooi, M., Mashile, K., Rambuda, T., Rangata, M.: A toolkit for text extraction and analysis for natural language processing tasks. In: *2022 International Conference on Artificial Intelligence, Big Data, Computing and Data Communication Systems (icABCD)*, pp. 1–6 (2022). <https://doi.org/10.1109/icABCD54961.2022.9856269>
741. Sehwag, V., Wang, S., Mittal, P., Jana, S.: Hydra: Pruning adversarially robust neural networks. *Advances in Neural Information Processing Systems* **33**, 19655–19666 (2020)
742. Serapio-García, G., Safdari, M., Crepy, C., Sun, L., Fitz, S., Romero, P., Abdulhai, M., Faust, A., Matarić, M.: Personality traits in large language models (2023)

743. Shahriari, B., Swersky, K., Wang, Z., Adams, R.P., De Freitas, N.: Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE* **104**(1), 148–175 (2015)
744. Shaik, T., Tao, X., Xie, H., Li, L., Zhu, X., Li, Q.: Exploring the landscape of machine unlearning: A survey and taxonomy. *arXiv preprint arXiv:2305.06360* **1**(2) (2023)
745. Shan, C., Weng, C., Wang, G., Su, D., Luo, M., Yu, D., Xie, L.: Component fusion: Learning replaceable language model component for end-to-end speech recognition system. In: *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5361–5635. IEEE (2019)
746. Shao, R., Shi, Z., Yi, J., Chen, P.Y., Hsieh, C.J.: On the adversarial robustness of vision transformers. *Transactions on Machine Learning Research* (2022)
747. Shao, W., Zhao, X., Ge, Y., Zhang, Z., Yang, L., Wang, X., Shan, Y., Luo, P.: Not all models are equal: predicting model transferability in a self-challenging fisher space. In: *European Conference on Computer Vision*, pp. 286–302. Springer (2022)
748. Shen, K., Leng, Y., Tan, X., Tang, S., Zhang, Y., Liu, W., Lin, E.: Mask the correct tokens: An embarrassingly simple approach for error correction. *arXiv preprint arXiv:2211.13252* (2022)
749. Shen, X., Chen, Z., Backes, M., Shen, Y., Zhang, Y.: “do anything now”: Characterizing and evaluating in-the-wild jailbreak prompts on large language models. *arXiv preprint arXiv:2308.03825* (2023)
750. Shetty, R., Schiele, B., Fritz, M.: {A4NT}: Author attribute anonymity by adversarial training of neural machine translation. In: *27th USENIX Security Symposium (USENIX Security 18)*, pp. 1633–1650 (2018)
751. Shi, F., Qing, P., Yang, D., Wang, N., Lei, Y., Lu, H., Lin, X.: Prompt space optimizing few-shot reasoning success with large language models. *arXiv preprint arXiv:2306.03799* (2023)
752. Shi, W., Han, X., Gonen, H., Holtzman, A., Tsvetkov, Y., Zettlemoyer, L.: Toward human readable prompt tuning: Kubrick’s the shining is a good movie, and a good prompt too? In: *EMNLP*, pp. 10994–11005 (2023)
753. Shi, Y., Daunhawer, I., Vogt, J.E., Torr, P.H.S., Sanyal, A.: How robust is unsupervised representation learning to distribution shift? (2022). <https://doi.org/10.48550/ARXIV.2206.08871>
754. Shi, Z., Wang, Y., Yin, F., Chen, X., Chang, K.W., Hsieh, C.J.: Red teaming language model detectors with language models. *arXiv preprint arXiv:2305.19713* (2023)
755. Shi, Z., Wei, J., Liang, Y.: A theoretical analysis on feature learning in neural networks: Emergence from inputs and advantage over fixed features. In: *International Conference on Learning Representations* (2021)
756. Shin, J., Lee, Y., Jung, K.: Effective sentence scoring method using bert for speech recognition. In: *Asian Conference on Machine Learning*, pp. 1081–1093. PMLR (2019)
757. Shin, R., Lin, C.H., Thomson, S., Chen, C., Roy, S., Platanios, E.A., Pauls, A., Klein, D., Eisner, J., Van Durme, B.: Constrained language models yield few-shot semantic parsers. In: *EMNLP* (2021)
758. Shin, T., Razeghi, Y., Logan IV, R.L., Wallace, E., Singh, S.: Autoprompt: Eliciting knowledge from language models with automatically generated prompts. *arXiv preprint arXiv:2010.15980* (2020)
759. Shrestha, P., Sierra, S., González, F.A., Montes, M., Rosso, P., Solorio, T.: Convolutional neural networks for authorship attribution of short texts. In: *Proceedings of the 15th conference of the European chapter of the association for computational linguistics: Volume 2, short papers*, pp. 669–674 (2017)
760. SIDDHARTHAN, A.: Ehud reiter and robert dale. building natural language generation systems. cambridge university press, 2000. *Natural Language Engineering* **7**(3), 271–274 (2001). <https://doi.org/10.1017/S1351324901212704>
761. Simmons, G.: Moral mimicry: Large language models produce moral rationalizations tailored to political identity. *arXiv preprint arXiv:2209.12106* (2022)

762. Singh, S.P., Alistarh, D.: Woodfisher: Efficient second-order approximation for neural network compression. *Advances in Neural Information Processing Systems* **33**, 18098–18109 (2020)
763. Sinha, A., Namkoong, H., Duchi, J.: Certifying some distributional robustness with principled adversarial training. In: *International Conference on Learning Representations* (2018)
764. Sivanandam, S., Deepa, S.: *Genetic Algorithms*. Springer Berlin Heidelberg (2008). https://doi.org/10.1007/978-3-540-73190-0_2
765. Slator: Large language models may leak personal data (2022). <https://slator.com/large-language-models-may-leak-personal-data/>
766. Snell, C., Zhong, R., Klein, D., Steinhardt, J.: Approximating how single head attention learns. *arXiv preprint arXiv:2103.07601* (2021)
767. Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C.D., Ng, A.Y., Potts, C.: Recursive deep models for semantic compositionality over a sentiment treebank. In: *Proceedings of the 2013 conference on empirical methods in natural language processing*, pp. 1631–1642 (2013)
768. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: *Proceedings of the International Conference on Machine Learning (ICML)* (2015)
769. Sohl-Dickstein, J., Weiss, E.A., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: *ICML* (2015)
770. Sohn, K., Chang, H., Lezama, J., Polania, L., Zhang, H., Hao, Y., Essa, I., Jiang, L.: Visual prompt tuning for generative transfer learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 19840–19851 (2023)
771. Sokol, K., Flach, P.: One explanation does not fit all: The promise of interactive explanations for machine learning transparency. *KI-Künstliche Intelligenz* **34**(2), 235–250 (2020)
772. Solaiman, I., Brundage, M., Clark, J., Askill, A., Herbert-Voss, A., Wu, J., Radford, A., Krueger, G., Kim, J.W., Kreps, S., et al.: Release strategies and the social impacts of language models. *arXiv preprint arXiv:1908.09203* (2019)
773. Solomon, D.H., Allen, K.D., Katz, P., Sawalha, A.H., Yelin, E.: Chatgpt, et al... artificial intelligence, authorship, and medical publishing. *ACR Open Rheumatology* **5**(6), 288 (2023)
774. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: *ICLR* (2021)
775. Song, Y., Durkan, C., Murray, I., Ermon, S.: Maximum likelihood training of score-based diffusion models. In: *NIPS* (2021)
776. Song, Y., Ermon, S.: Generative modeling by estimating gradients of the data distribution. In: *NIPS* (2019)
777. Song, Y., Ermon, S.: Improved techniques for training score-based generative models. In: *NIPS* (2020)
778. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: *ICLR* (2021)
779. Soomro, K., Zamir, A.R., Shah, M.: Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402* (2012)
780. Sridhara, G., G., R.H., Mazumdar, S.: (2023)
781. Sriram, A., Jun, H., Satheesh, S., Coates, A.: Cold fusion: Training seq2seq models together with language models. *arXiv preprint arXiv:1708.06426* (2017)
782. Staab, R., Vero, M., Balunović, M., Vechev, M.: Beyond memorization: Violating privacy via inference with large language models (2023)
783. Stoychev, S., Gunes, H.: The effect of model compression on fairness in facial expression recognition. *arXiv preprint arXiv:2201.01709* (2022)
784. Struppek, L., Hintersdorf, D., Kersting, K.: Rickrolling the artist: Injecting invisible backdoors into text-guided image generation models. In: *ArXiv* (2022)
785. Su, D., Zhang, H., Chen, H., Yi, J., Chen, P.Y., Gao, Y.: Is robustness the cost of accuracy?—a comprehensive study on the robustness of 18 deep image classification models. In: *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 631–648 (2018)

786. Su, Y., Collier, N.: Contrastive search is what you need for neural text generation. *Trans. Mach. Learn. Res.* **2023** (2023). URL <https://openreview.net/forum?id=GbkWw3jwL9>
787. Su, Y., Lan, T., Wang, Y., Yogatama, D., Kong, L., Collier, N.: A contrastive framework for neural text generation. In: *NeurIPS* (2022)
788. Sun, F., Liu, J., Wu, J., Pei, C., Lin, X., Ou, W., Jiang, P.: Bert4rec: Sequential recommendation with bidirectional encoder representations from transformer. In: *Proceedings of the 28th ACM international conference on information and knowledge management*, pp. 1441–1450 (2019)
789. Sun, H., Zhang, Z., Deng, J., Cheng, J., Huang, M.: Safety assessment of chinese large language models. *arXiv preprint arXiv:2304.10436* (2023)
790. Sun, H.L., Hsiung, L., Chandramoorthy, N., Chen, P.Y., Ho, T.Y.: NeuralFuse: Learning to Recover the Accuracy of Access-Limited Neural Network Inference in Low-Voltage Regimes. *arXiv preprint arXiv:2306.16869* (2023)
791. Sun, L., Huang, Y., Wang, H., Wu, S., Zhang, Q., Gao, C., Huang, Y., Lyu, W., Zhang, Y., Li, X., et al.: Trustllm: Trustworthiness in large language models. *International Conference on Machine Learning* (2024)
792. Sun, M., Liu, Z., Bair, A., Kolter, J.Z.: A simple and effective pruning approach for large language models. *arXiv preprint arXiv:2306.11695* (2023)
793. Sun, Z., Shen, Y., Zhou, Q., Zhang, H., Chen, Z., Cox, D., Yang, Y., Gan, C.: Principle-driven self-alignment of language models from scratch with minimal human supervision. *arXiv preprint arXiv:2305.03047* (2023)
794. Suresh, H., Gomez, S.R., Nam, K.K., Satyanarayan, A.: Beyond expertise and roles: A framework to characterize the stakeholders of interpretable machine learning and their needs. In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1–16 (2021)
795. Suresh, H., Guttig, J.: A framework for understanding sources of harm throughout the machine learning life cycle. In: *Equity and access in algorithms, mechanisms, and optimization*, pp. 1–9 (2021)
796. Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, R.: Intriguing properties of neural networks. *International Conference on Learning Representations* (2014)
797. Szymanski, P., Zelasko, P., Morzy, M., Szymczak, A., Zyla-Hoppe, M., Banaszczyk, J., Augustyniak, L., Mizgajski, J., Carmiel, Y.: Wer we are and wer we think we are. *arXiv preprint arXiv:2010.03432* (2020)
798. Takezawa, Y., Sato, R., Bao, H., Niwa, K., Yamada, M.: Necessary and sufficient watermark for large language models (2023)
799. Talat, Z., Blix, H., Valvoda, J., Ganesh, M.I., Cotterell, R., Williams, A.: A word on machine ethics: A response to jiang et al.(2021). *arXiv preprint arXiv:2111.04158* (2021)
800. Tang, Y., Han, K., Wang, Y., Xu, C., Guo, J., Xu, C., Tao, D.: Patch slimming for efficient vision transformers. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12165–12174 (2022)
801. Tang, Y., Qu, A., Chow, A.H., Lam, W.H., Wong, S., Ma, W.: Domain adversarial spatial-temporal network: a transferable framework for short-term traffic forecasting across cities. In: *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pp. 1905–1915 (2022)
802. Tang, Y.C., Chen, P.Y., Ho, T.Y.: Neural clamping: Joint input perturbation and temperature scaling for neural network calibration. *arXiv preprint arXiv:2209.11604* (2022)
803. Tang, Z., Rybin, D., Chang, T.H.: Zeroth-order optimization meets human feedback: Provable learning via ranking oracles. *arXiv preprint arXiv:2303.03751* (2023)
804. Taori, R., Gulrajani, I., Zhang, T., Dubois, Y., Li, X., Guestrin, C., Liang, P., Hashimoto, T.B.: Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca (2023)
805. Tarzanagh, D.A., Li, Y., Thrampoulidis, C., Oymak, S.: Transformers as support vector machines. *arXiv preprint arXiv:2308.16898* (2023)

806. Tarzanagh, D.A., Li, Y., Zhang, X., Oymak, S.: Max-margin token selection in attention mechanism. *CoRR* (2023)
807. Thirunavukarasu, A.J., Ting, D.S.J., Elangovan, K., Gutierrez, L., Tan, T.F., Ting, D.S.W.: Large language models in medicine. *Nature medicine* **29**(8), 1930–1940 (2023)
808. Thudi, A., Deza, G., Chandrasekaran, V., Papernot, N.: Unrolling sgd: Understanding factors influencing machine unlearning. *arXiv preprint arXiv:2109.13398* (2021)
809. Thudi, A., Deza, G., Chandrasekaran, V., Papernot, N.: Unrolling sgd: Understanding factors influencing machine unlearning. In: *2022 IEEE 7th European Symposium on Security and Privacy (EuroS&P)*, pp. 303–319. *IEEE* (2022)
810. Thynne, I., Goldring, J.: Accountability and control: Government officials and the exercise of power. (No Title) (1987)
811. Tian, Y., Gan, R., Song, Y., Zhang, J., Zhang, Y.: ChiMed-GPT: A Chinese Medical Large Language Model with Full Training Regime and Better Alignment to Human Preferences. *arXiv preprint arXiv:2311.06025* (2023)
812. Tian, Y., Wang, Y., Chen, B., Du, S.: Scan and snap: Understanding training dynamics and token composition in 1-layer transformer. *arXiv preprint arXiv:2305.16380* (2023)
813. Tian, Y., Wang, Y., Zhang, Z., Chen, B., Du, S.S.: Joma: Demystifying multilayer transformers via joint dynamics of mlp and attention. In: *Conference on Parsimony and Learning (Recent Spotlight Track)* (2023)
814. Topkara, U., Topkara, M., Atallah, M.J.: The hiding virtues of ambiguity: quantifiably resilient watermarking of natural language text through synonym substitutions. In: *MM&Sec*, pp. 164–174 (2006)
815. Torczon, V.: On the convergence of the multidirectional search algorithm. *SIAM journal on Optimization* **1**(1), 123–145 (1991)
816. Toshniwal, S., Kannan, A., Chiu, C.C., Wu, Y., Sainath, T.N., Livescu, K.: A comparison of techniques for language model integration in encoder-decoder speech recognition. In: *2018 IEEE spoken language technology workshop (SLT)*, pp. 369–375. *IEEE* (2018)
817. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jégou, H.: Training data-efficient image transformers & distillation through attention. In: *International Conference on Machine Learning*, pp. 10347–10357. *PMLR* (2021)
818. Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al.: Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023)
819. Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al.: Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023)
820. Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., Bikel, D., Blecher, L., Canton-Ferrer, C., Chen, M., Cucurull, G., Esiobu, D., Fernandes, J., Fu, J., Fu, W., Fuller, B., Gao, C., Goswami, V., Goyal, N., Hartshorn, A., Hosseini, S., Hou, R., Inan, H., Kardaş, M., Kerkez, V., Khabsa, M., Kloumann, I., Korenev, A., Koura, P.S., Lachaux, M., Lavril, T., Lee, J., Liskovich, D., Lu, Y., Mao, Y., Martinet, X., Mihaylov, T., Mishra, P., Molybog, I., Nie, Y., Poulton, A., Reizenstein, J., Rungta, R., Saladi, K., Schelten, A., Silva, R., Smith, E.M., Subramanian, R., Tan, X.E., Tang, B., Taylor, R., Williams, A., Kuan, J.X., Xu, P., Yan, Z., Zarov, I., Zhang, Y., Fan, A., Kambadur, M., Narang, S., Rodriguez, A., Stojnic, R., Edunov, S., Scialom, T.: Llama 2: Open foundation and fine-tuned chat models. *CoRR* **abs/2307.09288** (2023)
821. Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., Bikel, D., Blecher, L., Canton-Ferrer, C., Chen, M., Cucurull, G., Esiobu, D., Fernandes, J., Fu, J., Fu, W., Fuller, B., Gao, C., Goswami, V., Goyal, N., Hartshorn, A., Hosseini, S., Hou, R., Inan, H., Kardaş, M., Kerkez, V., Khabsa, M., Kloumann, I., Korenev, A., Koura, P.S., Lachaux, M., Lavril, T., Lee, J., Liskovich, D., Lu, Y., Mao, Y., Martinet, X., Mihaylov, T., Mishra, P., Molybog, I., Nie, Y., Poulton, A., Reizenstein, J., Rungta, R., Saladi, K., Schelten, A., Silva, R., Smith, E.M., Subramanian, R., Tan, X.E., Tang, B., Taylor, R., Williams, A., Kuan, J.X., Xu, P., Yan, Z., Zarov, I., Zhang,

- Y., Fan, A., Kambadur, M., Narang, S., Rodriguez, A., Stojnic, R., Edunov, S., Scialom, T.: Llama 2: Open foundation and fine-tuned chat models. CoRR **abs/2307.09288** (2023). <https://doi.org/10.48550/arXiv.2307.09288>
822. Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al.: Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288 (2023)
 823. Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al.: Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288 (2023)
 824. Tran, D., Liu, J., Dusenberry, M.W., Phan, D., Collier, M., Ren, J., Han, K., Wang, Z., Mariet, Z., Hu, H., et al.: Plex: Towards reliability using pretrained large model extensions. arXiv preprint arXiv:2207.07411 (2022)
 825. Trelis: flama 2 - function calling llama 2 (2023). URL <https://huggingface.co/Trelis/Llama-2-7b-chat-hf-function-calling>
 826. Tsai, Y.L., Hsu, C.Y., Xie, C., Lin, C.H., Chen, J.Y., Li, B., Chen, P.Y., Yu, C.M., Huang, C.Y.: Ring-a-bell! how reliable are concept removal methods for diffusion models? International Conference on Learning Representations (2024)
 827. Tsai, Y.Y., Chen, P.Y., Ho, T.Y.: Transfer learning without knowing: Reprogramming black-box machine learning models with scarce data and limited resources. In: International Conference on Machine Learning, pp. 9614–9624 (2020)
 828. Tsai, Y.Y., Mao, C., Yang, J.: Convolutional visual prompt for robust visual perception. In: Advances in Neural Information Processing Systems, vol. 36, pp. 27897–27921. Curran Associates, Inc. (2023)
 829. Tsao, H.A., Hsiung, L., Chen, P.Y., Liu, S., Ho, T.Y.: Autovp: An automated visual prompting framework and benchmark. International Conference on Learning Representations (2024)
 830. Tschandl, P., Rosendahl, C., Kittler, H.: The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. Scientific data **5**(1), 1–9 (2018)
 831. Tsimpoukelli, M., Menick, J.L., Cabi, S., Eslami, S., Vinyals, O., Hill, F.: Multimodal few-shot learning with frozen language models. Advances in Neural Information Processing Systems **34**, 200–212 (2021)
 832. Tu, C.C., Ting, P., Chen, P.Y., Liu, S., Zhang, H., Yi, J., Hsieh, C.J., Cheng, S.M.: Autozoom: Autoencoder-based zeroth order optimization method for attacking black-box neural networks. In: Proceedings of the AAAI Conference on Artificial Intelligence, pp. 742–749 (2019)
 833. Tu, T., Azizi, S., Driess, D., Schaeckermann, M., Amin, M., Chang, P.C., Carroll, A., Lau, C., Tanno, R., Ktena, I., et al.: Towards generalist biomedical ai. arXiv preprint arXiv:2307.14334 (2023)
 834. Tukey, J.W.: Mathematics and the picturing of data. In: Proceedings of the International Congress of Mathematicians, Vancouver, 1975, vol. 2, pp. 523–531 (1975)
 835. Tummala, S., Kadry, S., Bukhari, S.A.C., Rauf, H.T.: Classification of brain tumor from magnetic resonance imaging using vision transformers ensembling. Current Oncology **29**(10), 7498–7511 (2022)
 836. Turan, M.A.T., Vincent, E., Jouvet, D.: Achieving multi-accent asr via unsupervised acoustic model adaptation. In: INTERSPEECH 2020 (2020)
 837. Tüske, Z., Tahir, M.A., Schlüter, R., Ney, H.: Integrating gaussian mixtures into deep neural networks: Softmax layer with hidden variables. In: 2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 4285–4289. IEEE (2015)
 838. UBC: Reducing bias in llms (2023). <https://www.ischool.berkeley.edu/projects/2023/reducing-bias-large-language-models>
 839. Udagawa, T., Suzuki, M., Kurata, G., Itoh, N., Saon, G.: Effect and analysis of large-scale language model rescoring on competitive asr systems. arXiv preprint arXiv:2204.00212 (2022)

840. Ueoka, H., Murawaki, Y., Kurohashi, S.: Frustratingly easy edit-based linguistic steganography with a masked language model. In: NAACL, pp. 5486–5492 (2021)
841. Utpala, S., Hooker, S., Chen, P.Y.: Locally differentially private document generation using zero shot prompting. In: Findings of the Association for Computational Linguistics: EMNLP 2023, pp. 8442–8457. Association for Computational Linguistics (2023)
842. Variani, E., Chen, T., Apfel, J., Ramabhadran, B., Lee, S., Moreno, P.: Neural oracle search on n-best hypotheses. In: ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 7824–7828. IEEE (2020)
843. Varshney, N., Dolin, P., Seth, A., Baral, C.: The art of defending: A systematic evaluation and analysis of LLM defense strategies on safety and over-defensiveness. CoRR **abs/2401.00287** (2024)
844. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
845. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
846. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need (2023)
847. Veale, M., Zuiderveen Borgesius, F.: Demystifying the draft eu artificial intelligence act—analysing the good, the bad, and the unclear elements of the proposed approach. *Computer Law Review International* **22**(4), 97–112 (2021)
848. Voita, E., Titov, I.: Information-theoretic probing with minimum description length. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 183–196 (2020)
849. Von Oswald, J., Niklasson, E., Randazzo, E., Sacramento, J., Mordvintsev, A., Zhmoginov, A., Vladymyrov, M.: Transformers learn in-context by gradient descent. In: International Conference on Machine Learning, pp. 35151–35174. PMLR (2023)
850. Vu, T., Iyyer, M., Wang, X., Constant, N., Wei, J., Wei, J., Tar, C., Sung, Y.H., Zhou, D., Le, Q., et al.: Freshllms: Refreshing large language models with search engine augmentation. *arXiv preprint arXiv:2310.03214* (2023)
851. Wallach, W., Allen, C., Smit, I.: Machine morality: bottom-up and top-down approaches for modelling human moral faculties. *Ai & Society* **22**, 565–582 (2008)
852. Wan, Y., Pu, G., Sun, J., Garimella, A., Chang, K.W., Peng, N.: “kelly is a warm person, joseph is a role model”: Gender biases in llm-generated reference letters. *arXiv preprint arXiv:2310.09219* (2023)
853. Wang, A., Pruksachatkun, Y., Nangia, N., Singh, A., Michael, J., Hill, F., Levy, O., Bowman, S.R.: Superglue: A stickier benchmark for general-purpose language understanding systems (2020)
854. Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., Bowman, S.R.: Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461* (2018)
855. Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., Bowman, S.R.: GLUE: A multi-task benchmark and analysis platform for natural language understanding. In: International Conference on Learning Representations (2019)
856. Wang, B., Chen, W., Pei, H., Xie, C., Kang, M., Zhang, C., Xu, C., Xiong, Z., Dutta, R., Schaeffer, R., et al.: Decodingtrust: A comprehensive assessment of trustworthiness in gpt models. *arXiv preprint arXiv:2306.11698* (2023)
857. Wang, B., Chen, W., Pei, H., Xie, C., Kang, M., Zhang, C., Xu, C., Xiong, Z., Dutta, R., Schaeffer, R., et al.: Decodingtrust: A comprehensive assessment of trustworthiness in gpt models (2023)
858. Wang, B., Xu, C., Wang, S., Gan, Z., Cheng, Y., Gao, J., Awadallah, A.H., Li, B.: Adversarial glue: A multi-task benchmark for robustness evaluation of language models. *arXiv preprint arXiv:2111.02840* (2021)

859. Wang, C., Cheng, S., Xu, Z., Ding, B., Wang, Y., Zhang, Y.: Evaluating open question answering evaluation. *arXiv preprint arXiv:2305.12421* (2023)
860. Wang, G., Cheng, S., Zhan, X., Li, X., Song, S., Liu, Y.: Openchat: Advancing open-source language models with mixed-quality data. *arXiv preprint arXiv:2309.11235* (2023)
861. Wang, H., Fu, T., Du, Y., Gao, W., Huang, K., Liu, Z., Chandak, P., Liu, S., Van Katwyk, P., Deac, A., et al.: Scientific discovery in the age of artificial intelligence. *Nature* **620**(7972), 47–60 (2023)
862. Wang, J., Hu, X., Hou, W., Chen, H., Zheng, R., Wang, Y., Yang, L., Huang, H., Ye, W., Geng, X., Jiao, B., Zhang, Y., Xie, X.: On the robustness of chatgpt: An adversarial and out-of-distribution perspective (2023)
863. Wang, J., Lan, C., Liu, C., Ouyang, Y., Qin, T., Lu, W., Chen, Y., Zeng, W., Yu, P.S.: Generalizing to unseen domains: A survey on domain generalization (2022)
864. Wang, L., Song, M., Rezapour, R., Kwon, B.C., Huh-Yoo, J.: People’s perceptions toward bias and related concepts in large language models: A systematic review. *arXiv preprint arXiv:2309.14504* (2023)
865. Wang, L., Yang, W., Chen, D., Zhou, H., Lin, Y., Meng, F., Zhou, J., Sun, X.: Towards codable text watermarking for large language models. *CoRR abs/2307.15992* (2023). <https://doi.org/10.48550/arXiv.2307.15992>
866. Wang, L., Yang, W., Chen, D., Zhou, H., Lin, Y., Meng, F., Zhou, J., Sun, X.: Towards codable watermarking for injecting multi-bit information to llm (2023)
867. Wang, P., Li, L., Chen, L., Zhu, D., Lin, B., Cao, Y., Liu, Q., Liu, T., Sui, Z.: Large language models are not fair evaluators. *arXiv preprint arXiv:2305.17926* (2023)
868. Wang, W., Hu, K., Sainath, T.N.: Deliberation of streaming rnn-transducer by non-autoregressive decoding. In: *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7452–7456. IEEE (2022)
869. Wang, W., Tu, Z., Chen, C., Yuan, Y., tse Huang, J., Jiao, W., Lyu, M.R.: All languages matter: On the multilingual safety of large language models (2023)
870. Wang, W., Xie, E., Li, X., Fan, D.P., Song, K., Liang, D., Lu, T., Luo, P., Shao, L.: Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 568–578 (2021)
871. Wang, X., Guo, W., Su, J., Yang, X., Yan, J.: Zarts: On zero-order optimization for neural architecture search. *Advances in Neural Information Processing Systems* **35**, 12868–12880 (2022)
872. Wang, Y., Chu, Z., Ouyang, X., Wang, S., Hao, H., Shen, Y., Gu, J., Xue, S., Zhang, J.Y., Cui, Q., et al.: Enhancing recommender systems with large language model reasoning graphs. *arXiv preprint arXiv:2308.10835* (2023)
873. Wang, Y., Kordi, Y., Mishra, S., Liu, A., Smith, N.A., Khoshabi, D., Hajishirzi, H.: Self-instruct: Aligning language model with self generated instructions. *arXiv preprint arXiv:2212.10560* (2022)
874. Wang, Y., Li, H., Han, X., Nakov, P., Baldwin, T.: Do-not-answer: A dataset for evaluating safeguards in llms. *arXiv preprint arXiv:2308.13387* (2023)
875. Wang, Y., Wang, Y., Dang, K., Liu, J., Liu, Z.: A comprehensive survey of grammatical error correction. *ACM Transactions on Intelligent Systems and Technology (TIST)* **12**(5), 1–51 (2021)
876. Wang, Z., Hunt, J.J., Zhou, M.: Diffusion policies as an expressive policy class for offline reinforcement learning. In: *CoRR* (2022)
877. Wang, Z., Jiang, W., Zhu, Y.M., Yuan, L., Song, Y., Liu, W.: Dynamixer: a vision mlp architecture with dynamic mixing. In: *International Conference on Machine Learning*, pp. 22691–22701. PMLR (2022)
878. Wang, Z., Jiang, Y., Lu, Y., He, P., Chen, W., Wang, Z., Zhou, M., et al.: In-context learning unlocked for diffusion models. *Advances in Neural Information Processing Systems* **36**, 8542–8562 (2023)

879. Wang, Z., Le, T., Lee, D.: Upton: Unattributable authorship text via data poisoning. arXiv preprint arXiv:2211.09717 (2022)
880. Wang, Z., Li, R., Dong, B., Wang, J., Li, X., Liu, N., Mao, C., Zhang, W., Dong, L., Gao, J., Wang, J.: Can llms like gpt-4 outperform traditional ai tools in dementia diagnosis? maybe, but not today (2023)
881. Warnecke, A., Pirch, L., Wressnegger, C., Rieck, K.: Machine unlearning of features and labels. arXiv preprint arXiv:2108.11577 (2021)
882. Wei, A., Haghtalab, N., Steinhardt, J.: Jailbroken: How does llm safety training fail? arXiv preprint arXiv:2307.02483 (2023)
883. Wei, A., Haghtalab, N., Steinhardt, J.: Jailbroken: How does LLM safety training fail? CoRR **abs/2307.02483** (2023)
884. Wei, C., Chen, Y., Ma, T.: Statistically meaningful approximation: a case study on approximating turing machines with transformers. arXiv preprint arXiv:2107.13163 (2021)
885. Wei, J., Bosma, M., Zhao, V., Guu, K., Yu, A.W., Lester, B., Du, N., Dai, A.M., Le, Q.V.: Finetuned language models are zero-shot learners. In: International Conference on Learning Representations (2022)
886. Wei, J., Bosma, M., Zhao, V.Y., Guu, K., Yu, A.W., Lester, B., Du, N., Dai, A.M., Le, Q.V.: Finetuned language models are zero-shot learners. arXiv preprint arXiv:2109.01652 (2021)
887. Wei, J., Tay, Y., Bommasani, R., Raffel, C., Zoph, B., Borgeaud, S., Yogatama, D., Bosma, M., Zhou, D., Metzler, D., et al.: Emergent abilities of large language models. arXiv preprint arXiv:2206.07682 (2022)
888. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q., Zhou, D.: Chain-of-thought prompting elicits reasoning in large language models (2023)
889. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q.V., Zhou, D., et al.: Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems* **35**, 24824–24837 (2022)
890. Wei, T., Luan, J., Liu, W., Dong, S., Wang, B.: Cmath: Can your language model pass chinese elementary school math test? (2023)
891. Wei, Z., Wang, Y., Wang, Y.: Jailbreak and guard aligned language models with only few in-context demonstrations (2023)
892. Wei, Z., Wang, Y., Wang, Y.: Jailbreak and guard aligned language models with only few in-context demonstrations. CoRR **abs/2310.06387** (2023)
893. Welbl, J., Glaese, A., Uesato, J., Dathathri, S., Mellor, J., Hendricks, L.A., Anderson, K., Kohli, P., Coppin, B., Huang, P.S.: Challenges in detoxifying language models. arXiv preprint arXiv:2109.07445 (2021)
894. Welleck, S., Kulikov, I., Roller, S., Dinan, E., Cho, K., Weston, J.: Neural text generation with unlikelihood training. In: 8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020 (2020)
895. Wen, Q., Zhou, T., Zhang, C., Chen, W., Ma, Z., Yan, J., Sun, L.: Transformers in time series: A survey. In: International Joint Conference on Artificial Intelligence (2023)
896. Wen, W., Wu, C., Wang, Y., Chen, Y., Li, H.: Learning structured sparsity in deep neural networks. *Advances in neural information processing systems* **29** (2016)
897. Wen, Y., Jain, N., Kirchenbauer, J., Goldblum, M., Geiping, J., Goldstein, T.: Hard prompts made easy: Gradient-based discrete optimization for prompt tuning and discovery. In: *NeurIPS*, vol. 36 (2024)
898. Wen, Z., Li, Y.: Toward understanding the feature learning process of self-supervised contrastive learning. In: International Conference on Machine Learning, pp. 11112–11122. PMLR (2021)
899. Whitney, W.F., Song, M.J., Brandfonbrener, D., Altosaar, J., Cho, K.: Evaluating representations by the complexity of learning low-loss predictors. arXiv preprint arXiv:2009.07368 (2020)
900. Wichers, N., Denison, C., Beirami, A.: Gradient-based language model red teaming. In: *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)* (2024)

901. Wired: How chatgpt works: A look inside large language models (2023). URL <https://www.wired.com/story/how-chatgpt-works-large-language-model/>
902. Wirth, J., Peinl, R.: Automatic speech recognition in german: A detailed error analysis. In: 2022 IEEE International Conference on Omni-layer Intelligent Systems (COINS), pp. 1–8. IEEE (2022)
903. Wong, E., Santurkar, S., Madry, A.: Leveraging sparse linear layers for debuggable deep networks. In: International Conference on Machine Learning, pp. 11205–11216. PMLR (2021)
904. Woo, G., Liu, C., Sahoo, D., Kumar, A., Hoi, S.: Etsformer: Exponential smoothing transformers for time-series forecasting. arXiv preprint arXiv:2202.01381 (2022)
905. Wortsman, M., Ilharco, G., Gadre, S.Y., Roelofs, R., Gontijo-Lopes, R., Morcos, A.S., Namkoong, H., Farhadi, A., Carmon, Y., Kornblith, S., et al.: Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In: International Conference on Machine Learning, pp. 23965–23998. PMLR (2022)
906. Wu, H., Hu, T., Liu, Y., Zhou, H., Wang, J., Long, M.: Timesnet: Temporal 2d-variation modeling for general time series analysis. In: International Conference on Learning Representations (2023)
907. Wu, H., Xu, J., Wang, J., Long, M.: Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. *Advances in Neural Information Processing Systems* **34**, 22419–22430 (2021)
908. Wu, J., Li, X., Wei, C., Wang, H., Yuille, A., Zhou, Y., Xie, C.: Unleashing the power of visual prompting at the pixel level. arXiv preprint arXiv:2212.10556 (2022)
909. Wu, J., Zou, D., Chen, Z., Braverman, V., Gu, Q., Bartlett, P.L.: How many pretraining tasks are needed for in-context learning of linear regression? arXiv preprint arXiv:2310.08391 (2023)
910. Wu, Q., Li, L., Yu, Z.: Textgail: Generative adversarial imitation learning for text generation. In: AAAI, pp. 14067–14075 (2021)
911. Wu, S., Irsoy, O., Lu, S., Dabrowski, V., Dredze, M., Gehrmann, S., Kambadur, P., Rosenberg, D., Mann, G.: Bloomberggpt: A large language model for finance (2023)
912. Wu, T., Terry, M., Cai, C.J.: Ai chains: Transparent and controllable human-ai interaction by chaining large language model prompts. In: Proceedings of the 2022 CHI conference on human factors in computing systems, pp. 1–22 (2022)
913. Wu, X., Fu, X., Liu, Y., Lim, E.P., Hoi, S.C., Sun, Q.: A large-scale benchmark for food image segmentation. In: Proceedings of the 29th ACM International Conference on Multimedia, pp. 506–515 (2021)
914. Wu, Y., Hu, Z., Zhang, H., Huang, H.: Dipmark: A stealthy, efficient and resilient watermark for large language models (2023)
915. Wu, Z., Wang, Y., Ye, J., Kong, L.: Self-adaptive in-context learning: An information compression perspective for in-context example selection and ordering. *ACL* (2023)
916. Xia, Y., Tian, F., Wu, L., Lin, J., Qin, T., Yu, N., Liu, T.Y.: Deliberation networks: Sequence generation beyond one-pass decoding. *Advances in neural information processing systems* **30** (2017)
917. Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems* **34**, 12077–12090 (2021)
918. Xie, Q., Han, W., Zhang, X., Lai, Y., Peng, M., Lopez-Lira, A., Huang, J.: Pixiu: A large language model, instruction data and evaluation benchmark for finance (2023)
919. Xie, Q., Luong, M.T., Hovy, E., Le, Q.V.: Self-training with noisy student improves imagenet classification. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 10687–10698 (2020)
920. Xie, Y., Yi, J., Shao, J., Curl, J., Lyu, L., Chen, Q., Xie, X., Wu, F.: Defending chatgpt against jailbreak attack via self-reminders. *Nat. Mac. Intell.* **5**(12), 1486–1496 (2023)
921. Xiong, C., Qi, X., Chen, P.Y., Ho, T.Y.: Defensive prompt patch: A robust and interpretable defense of llms against jailbreak attacks. arXiv preprint arXiv:2405.20099 (2024)

922. Xiong, X., Liu, S., Li, D., Cai, Z., Niu, X.: A comprehensive survey on local differential privacy. *Security and Communication Networks* **2020**, 1–29 (2020)
923. Xu, G., Liu, J., Yan, M., Xu, H., Si, J., Zhou, Z., Yi, P., Gao, X., Sang, J., Zhang, R., Zhang, J., Peng, C., Huang, F., Zhou, J.: Cvalues: Measuring the values of chinese large language models from safety to responsibility (2023)
924. Xu, G., Liu, J., Yan, M., Xu, H., Si, J., Zhou, Z., Yi, P., Gao, X., Sang, J., Zhang, R., et al.: Cvalues: Measuring the values of chinese large language models from safety to responsibility. *arXiv preprint arXiv:2307.09705* (2023)
925. Xu, H., Zhu, T., Zhang, L., Zhou, W., Yu, P.S.: Machine unlearning: A survey. *ACM Computing Surveys* **56**(1), 1–36 (2023)
926. Xu, J., Ren, X., Lin, J., Sun, X.: Diversity-promoting GAN: A cross-entropy based generative adversarial network for diversified text generation. In: *EMNLP*, pp. 3940–3949 (2018)
927. Xu, L., Zhao, K., Zhu, L., Xue, H.: Sc-safety: A multi-round open-ended question adversarial safety benchmark for large language models in chinese. *arXiv preprint arXiv:2310.05818* (2023)
928. Xu, M.: Medicalgpt: Training medical gpt model. <https://github.com/shibing624/MedicalGPT> (2023)
929. Xu, X., Kong, K., Liu, N., Cui, L., Wang, D., Zhang, J., Kankanhalli, M.: An llm can fool itself: A prompt-based adversarial attack. *arXiv preprint arXiv:2310.13345* (2023)
930. Xu, Y., Zhao, S., Song, J., Stewart, R., Ermon, S.: A theory of usable information under computational constraints. In: *International Conference on Learning Representations* (2019)
931. Xu, Z., Aggarwal, A., Feyisetan, O., Teissier, N.: A differentially private text perturbation method using regularized mahalanobis metric. In: *Proceedings of the Second Workshop on Privacy in NLP*, pp. 7–17 (2020)
932. Xue, H., Salim, F.D.: Prompt-based time series forecasting: A new task and dataset. *arXiv preprint arXiv:2210.08964* (2022)
933. Xu, J., Wang, Y.C., Wei, C., Liu, X., Woo, J., Kuo, C.C.J.: Bias and fairness in chatbots: An overview. *arXiv preprint arXiv:2309.08836* (2023)
934. Yang, C.H.H., Gu, Y., Liu, Y.C., Ghosh, S., Bulyko, I., Stolcke, A.: Generative speech recognition error correction with large language models and task-activating prompting. In: *Proc. IEEE ASRU* (2023)
935. Yang, C.H.H., Li, B., Zhang, Y., Chen, N., Prabhavalkar, R., Sainath, T.N., Strohmaier, T.: From english to more languages: Parameter-efficient model reprogramming for cross-lingual speech recognition. In: *Proc. ICASSP*, pp. 1–5. *IEEE* (2023)
936. Yang, C.H.H., Tsai, Y.Y., Chen, P.Y.: Voice2series: Reprogramming acoustic models for time series classification. In: *International Conference on Machine Learning* (2021)
937. Yang, G.: Tensor programs ii: Neural tangent kernel for any architecture. *arXiv preprint arXiv:2006.14548* (2020)
938. Yang, H., Liang, Y., Guo, X., Wu, L., Wang, Z.: Theoretical characterization of how neural network pruning affects its generalization. *arXiv preprint arXiv:2301.00335* (2023)
939. Yang, H., Wang, Z.: On the neural tangent kernel analysis of randomly pruned neural networks. In: *International Conference on Artificial Intelligence and Statistics*, pp. 1513–1553. *PMLR* (2023)
940. Yang, K.C., Menczer, F.: Large language models can rate news outlet credibility (2023)
941. Yang, L., Zhang, S., Qin, L., Li, Y., Wang, Y., Liu, H., Wang, J., Xie, X., Zhang, Y.: Glue-x: Evaluating natural language understanding models from an out-of-distribution generalization perspective. *arXiv preprint arXiv:2211.08073* (2022)
942. Yang, X., Gao, J., Xue, W., Alexandersson, E.: Pllama: An open-source large language model for plant science (2024)
943. Yang, X., Wang, X., Zhang, Q., Petzold, L., Wang, W.Y., Zhao, X., Lin, D.: Shadow alignment: The ease of subverting safely-aligned language models (2023)
944. Yang, Z., Hu, Z., Dyer, C., Xing, E.P., Berg-Kirkpatrick, T.: Unsupervised text style transfer using language models as discriminators. In: *NeurIPS*, pp. 7298–7309 (2018)

945. Yang, Z., Sha, Z., Backes, M., Zhang, Y.: From visual prompt learning to zero-shot transfer: Mapping is all you need. arXiv preprint arXiv:2303.05266 (2023)
946. Yao, A.C.C.: How to generate and exchange secrets. In: 27th Annual Symposium on Foundations of Computer Science (sfcs 1986), pp. 162–167 (1986). <https://doi.org/10.1109/SFCS.1986.25>
947. Yao, D., Zhang, J., Harris, I.G., Carlsson, M.: Fuzzllm: A novel and universal fuzzing framework for proactively discovering jailbreak vulnerabilities in large language models (2023)
948. Yao, Y., Xu, X., Liu, Y.: Large language model unlearning. arXiv preprint arXiv:2310.10683 (2023)
949. Ye, H., Huang, Z., Fang, C., Li, C.J., Zhang, T.: Hessian-aware zeroth-order optimization for black-box adversarial attack. arXiv preprint arXiv:1812.11377 (2018)
950. Ye, W., Ou, M., Li, T., Ma, X., Yanggong, Y., Wu, S., Fu, J., Chen, G., Zhao, J., et al.: Assessing hidden risks of llms: An empirical study on robustness, consistency, and credibility. arXiv preprint arXiv:2305.10235 (2023)
951. Yen, H., Ku, P.J., Yang, C.H.H., Hu, H., Siniscalchi, S.M., Chen, P.Y., Tsao, Y.: Neural model reprogramming with similarity based mapping for low-resource spoken command classification. arXiv preprint arXiv:2110.03894 (2021)
952. Yin, H., Vahdat, A., Alvarez, J.M., Mallya, A., Kautz, J., Molchanov, P.: A-vit: Adaptive tokens for efficient vision transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 10809–10818 (2022)
953. Yin, S., Fu, C., Zhao, S., Li, K., Sun, X., Xu, T., Chen, E.: A survey on multimodal large language models. arXiv preprint arXiv:2306.13549 (2023)
954. Yin, Z., Sun, Q., Guo, Q., Wu, J., Qiu, X., Huang, X.: Do large language models know what they don't know? arXiv preprint arXiv:2305.18153 (2023)
955. Yong, Z.X., Menghini, C., Bach, S.H.: Low-resource languages jailbreak gpt-4 (2023)
956. Yong, Z.X., Menghini, C., Bach, S.H.: Low-resource languages jailbreak GPT-4. CoRR **abs/2310.02446** (2023)
957. Yong, Z.X., Menghini, C., Bach, S.H.: Low-resource languages jailbreak gpt-4 (2024)
958. Yoo, K., Ahn, W., Kwak, N.: Advancing beyond identification: Multi-bit watermark for large language models (2023)
959. You, K., Liu, Y., Wang, J., Long, M.: Logme: Practical assessment of pre-trained models for federated learning. In: International Conference on Machine Learning, pp. 12133–12143. PMLR (2021)
960. Yu, C., Jeoung, S., Kasi, A., Yu, P., Ji, H.: Unlearning bias in language models by partitioning gradients. In: Findings of the Association for Computational Linguistics: ACL 2023, pp. 6032–6048 (2023)
961. Yu, D., Naik, S., Backurs, A., Gopi, S., Inan, H.A., Kamath, G., Kulkarni, J., Lee, Y.T., Manoel, A., Wutschitz, L., et al.: Differentially private fine-tuning of language models. In: International Conference on Learning Representations (2021)
962. Yu, J., Lin, X., Yu, Z., Xing, X.: Gptfuzzer: Red teaming large language models with auto-generated jailbreak prompts (2023)
963. Yu, J., Wang, Z., Vasudevan, V., Yeung, L., Seyedhosseini, M., Wu, Y.: Coca: Contrastive captioners are image-text foundation models. arXiv preprint arXiv:2205.01917 (2022)
964. Yu, J., Zhu, J., Wang, Y., Liu, Y., Chang, H., Nie, J., Kong, C., Chong, R., XinLiu, An, J., Lu, L., Fang, M., Zhu, L.: Taoli llama. <https://github.com/blucicall/taoli> (2023)
965. Yu, L., Chen, Q., Lin, J., He, L.: Black-box prompt tuning for vision-language model as a service. In: IJCAI, pp. 1686–1694 (2023)
966. Yu, L., Zhang, W., Wang, J., Yu, Y.: Seqgan: Sequence generative adversarial nets with policy gradient. In: AAAI, pp. 2852–2858 (2017)
967. Yuan, A., Coenen, A., Reif, E., Ippolito, D.: Wordcraft: story writing with large language models. In: 27th International Conference on Intelligent User Interfaces, pp. 841–852 (2022)
968. Yuan, Z., Xue, H., Wang, X., Liu, Y., Zhao, Z., Wang, K.: Artgpt-4: Artistic vision-language understanding with adapter-enhanced minigpt-4 (2023)

969. Yuan, Z., Yuan, H., Tan, C., Wang, W., Huang, S.: How well do large language models perform in arithmetic tasks? (2023)
970. Yue, S., Chen, W., Wang, S., Li, B., Shen, C., Liu, S., Zhou, Y., Xiao, Y., Yun, S., Huang, X., Wei, Z.: Disc-lawllm: Fine-tuning large language models for intelligent legal services (2023)
971. Yun, C., Bhojanapalli, S., Rawat, A.S., Reddi, S., Kumar, S.: Are transformers universal approximators of sequence-to-sequence functions? In: International Conference on Learning Representations (2019)
972. Zaken, E.B., Ravfogel, S., Goldberg, Y.: Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. arXiv preprint arXiv:2106.10199 (2021)
973. Zellers, R., Holtzman, A., Rashkin, H., Bisk, Y., Farhadi, A., Roesner, F., Choi, Y.: Defending against neural fake news. In: NeurIPS, pp. 9051–9062 (2019)
974. Zeng, A., Chen, M., Zhang, L., Xu, Q.: Are transformers effective for time series forecasting? In: Proceedings of the AAAI conference on artificial intelligence, vol. 37, pp. 11121–11128 (2023)
975. Zeyer, A., Irie, K., Schlüter, R., Ney, H.: Improved training of end-to-end attention models for speech recognition. arXiv preprint arXiv:1805.03294 (2018)
976. Zhang, B., Haddow, B., Birch, A.: Prompting large language model for machine translation: A case study. In: A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, J. Scarlett (eds.) International Conference on Machine Learning, ICML 2023, 23–29 July 2023, Honolulu, Hawaii, USA, *Proceedings of Machine Learning Research*, vol. 202, pp. 41092–41110. PMLR (2023). URL <https://proceedings.mlr.press/v202/zhang23m.html>
977. Zhang, C., Bengio, S., Hardt, M., Recht, B., Vinyals, O.: Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM* **64**(3), 107–115 (2021)
978. Zhang, E., Wang, K., Xu, X., Wang, Z., Shi, H.: Forget-me-not: Learning to forget in text-to-image diffusion models. arXiv preprint arXiv:2303.17591 (2023)
979. Zhang, E., Wang, K., Xu, X., Wang, Z., Shi, H.: Forget-me-not: Learning to forget in text-to-image diffusion models. arXiv preprint arXiv:2303.17591 (2023)
980. Zhang, H., Chen, J., Jiang, F., Yu, F., Chen, Z., Li, J., Chen, G., Wu, X., Zhang, Z., Xiao, Q., Wan, X., Wang, B., Li, H.: Huatuogpt, towards taming language models to be a doctor. arXiv preprint arXiv:2305.15075 (2023)
981. Zhang, J., Chen, S., Liu, J., He, J.: Composing parameter-efficient modules with arithmetic operations. arXiv preprint arXiv:2306.14870 (2023)
982. Zhang, K., Wen, Q., Zhang, C., Cai, R., Jin, M., Liu, Y., Zhang, J., Liang, Y., Pang, G., Song, D., et al.: Self-supervised learning for time series analysis: Taxonomy, progress, and prospects. arXiv preprint arXiv:2306.10125 (2023)
983. Zhang, K., Yu, J., Yan, Z., Liu, Y., Adhikarla, E., Fu, S., Chen, X., Chen, C., Zhou, Y., Li, X., He, L., Davison, B.D., Li, Q., Chen, Y., Liu, H., Sun, L.: Biomedgpt: A unified and generalist biomedical generative pre-trained transformer for vision, language, and multimodal tasks (2023)
984. Zhang, L., Cai, W., Liu, Z., Yang, Z., Dai, W., Liao, Y., Qin, Q., Li, Y., Liu, X., Liu, Z., Zhu, Z., Wu, A., Guo, X., Chen, Y.: Fineval: A chinese financial domain knowledge evaluation benchmark for large language models (2023)
985. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 3836–3847 (2023)
986. Zhang, M., Ré, C.: Contrastive adapters for foundation model group robustness. arXiv preprint arXiv:2207.07180 (2022)
987. Zhang, Q., Chen, Y.: Fast sampling of diffusion models with exponential integrator. In: ICLR (2023)
988. Zhang, R., Frei, S., Bartlett, P.L.: Trained transformers learn linear models in-context. arXiv preprint arXiv:2306.09927 (2023)

989. Zhang, R., Han, J., Zhou, A., Hu, X., Yan, S., Lu, P., Li, H., Gao, P., Qiao, Y.: Llama-adapter: Efficient fine-tuning of language models with zero-init attention. arXiv preprint arXiv:2303.16199 (2023)
990. Zhang, R., Lin, L., Bai, Y., Mei, S.: Negative preference optimization: From catastrophic collapse to effective unlearning. arXiv preprint arXiv:2404.05868 (2024)
991. Zhang, R., Wang, Y.S., Yang, Y.: Generation-driven contrastive self-training for zero-shot text classification with instruction-tuned gpt (2023)
992. Zhang, S., Lei, M., Yan, Z.: Automatic spelling correction with transformer for ctc-based end-to-end speech recognition. arXiv preprint arXiv:1904.10045 (2019)
993. Zhang, S., Roller, S., Goyal, N., Artetxe, M., Chen, M., Chen, S., Dewan, C., Diab, M., Li, X., Lin, X.V., et al.: Opt: Open pre-trained transformer language models. arXiv preprint arXiv:2205.01068 (2022)
994. Zhang, S., Wang, M., Chen, P.Y., Liu, S., Lu, S., Liu, M.: Joint edge-model sparse learning is provably efficient for graph neural networks. In: The Eleventh International Conference on Learning Representations (2023)
995. Zhang, S., Wang, M., Liu, S., Chen, P.Y., Xiong, J.: Fast learning of graph neural networks with guaranteed generalizability: One-hidden-layer case. In: International Conference on Machine Learning, pp. 11268–11277. PMLR (2020)
996. Zhang, S., Wang, M., Liu, S., Chen, P.Y., Xiong, J.: Why lottery ticket wins? a theoretical perspective of sample complexity on sparse neural networks. *Advances in Neural Information Processing Systems* **34** (2021)
997. Zhang, S., Wang, M., Xiong, J., Liu, S., Chen, P.Y.: Improved linear convergence of training cnns with generalizability guarantees: A one-hidden-layer case. *IEEE Transactions on Neural Networks and Learning Systems* (2020)
998. Zhang, T., Ladhak, F., Durmus, E., Liang, P., McKeown, K., Hashimoto, T.B.: Benchmarking large language models for news summarization (2023)
999. Zhang, T., Zhang, Y., Cao, W., Bian, J., Yi, X., Zheng, S., Li, J.: Less is more: Fast multivariate time series forecasting with light sampling-oriented mlp structures. arXiv preprint arXiv:2207.01186 (2022)
1000. Zhang, W., Aljunied, S.M., Gao, C., Chia, Y.K., Bing, L.: M3exam: A multilingual, multimodal, multilevel benchmark for examining large language models (2023)
1001. Zhang, W., Deng, Y., Liu, B., Pan, S.J., Bing, L.: Sentiment analysis in the era of large language models: A reality check (2023)
1002. Zhang, X., Li, C., Zong, Y., Ying, Z., He, L., Qiu, X.: Evaluating the performance of large language models on gaokao benchmark (2023)
1003. Zhang, X., Tian, C., Yang, X., Chen, L., Li, Z., Petzold, L.R.: Alpacre:instruction-tuned large language models for medical application (2023)
1004. Zhang, X., Wang, L., Helwig, J., Luo, Y., Fu, C., Xie, Y., Liu, M., Lin, Y., Xu, Z., Yan, K., Adams, K., Weiler, M., Li, X., Fu, T., Wang, Y., Yu, H., Xie, Y., Fu, X., Strasser, A., Xu, S., Liu, Y., Du, Y., Saxton, A., Ling, H., Lawrence, H., Stärk, H., Gui, S., Edwards, C., Gao, N., Ladera, A., Wu, T., Hofgard, E.F., Tehrani, A.M., Wang, R., Daigavane, A., Bohde, M., Kurtin, J., Huang, Q., Phung, T., Xu, M., Joshi, C.K., Mathis, S.V., Azizzadenesheli, K., Fang, A., Aspuru-Guzik, A., Bekkers, E., Bronstein, M., Zitnik, M., Anandkumar, A., Ermon, S., Liò, P., Yu, R., Günnemann, S., Leskovec, J., Ji, H., Sun, J., Barzilay, R., Jaakkola, T., Coley, C.W., Qian, X., Qian, X., Smidt, T., Ji, S.: Artificial intelligence for science in quantum, atomistic, and continuum systems. arXiv preprint arXiv:2307.08423 (2023)
1005. Zhang, X., Zhao, Z., Tsiligkaridis, T., Zitnik, M.: Self-supervised contrastive pre-training for time series via time-frequency consistency. *Advances in Neural Information Processing Systems* **35**, 3988–4003 (2022)
1006. Zhang, Y., Ding, L., Zhang, L., Tao, D.: Intention analysis prompting makes large language models A good jailbreak defender. *CoRR* **abs/2401.06561** (2024)
1007. Zhang, Y., Jia, J., Chen, X., Chen, A., Zhang, Y., Liu, J., Ding, K., Liu, S.: To generate or not? safety-driven unlearned diffusion models are still easy to generate unsafe images... for now. *European Conference on Computer Vision (ECCV)* (2024)

1008. Zhang, Y., Khanduri, P., Tsaknakis, I., Yao, Y., Hong, M., Liu, S.: An introduction to bilevel optimization: Foundations and applications in signal processing and machine learning. *IEEE Signal Processing Magazine* **41**(1), 38–59 (2024)
1009. Zhang, Y., Li, P., Hong, J., Li, J., Zhang, Y., Zheng, W., Chen, P.Y., Lee, J.D., Yin, W., Hong, M., et al.: Revisiting zeroth-order optimization for memory-efficient llm fine-tuning: A benchmark. In: *International Conference on Machine Learning* (2024)
1010. Zhang, Y., Yao, Y., Jia, J., Yi, J., Hong, M., Chang, S., Liu, S.: How to robustify black-box ML models? a zeroth-order optimization perspective. In: *International Conference on Learning Representations* (2022)
1011. Zhang, Y., Yao, Y., Ram, P., Zhao, P., Chen, T., Hong, M., Wang, Y., Liu, S.: Advancing model pruning via bi-level optimization. In: *Advances in Neural Information Processing Systems* (2022)
1012. Zhang, Y., Zhang, Y., Yao, Y., Jia, J., Liu, J., Liu, X., Liu, S.: Unlearncanvas: A stylized image dataset to benchmark machine unlearning for diffusion models. *arXiv preprint arXiv:2402.11846* (2024)
1013. Zhang, Z., Lei, L., Wu, L., Sun, R., Huang, Y., Long, C., Liu, X., Lei, X., Tang, J., Huang, M.: Safetybench: Evaluating the safety of large language models with multiple choice questions (2023)
1014. Zhang, Z., Yang, J., Ke, P., Huang, M.: Defending large language models against jailbreaking attacks through goal prioritization (2023)
1015. Zhao, P., Liu, S., Chen, P.Y., Hoang, N., Xu, K., Kailkhura, B., Lin, X.: On the design of black-box adversarial examples by leveraging gradient-free optimization and operator splitting method. In: *IEEE International Conference on Computer Vision*, pp. 121–130 (2019)
1016. Zhao, W., Bai, L., Rao, Y., Zhou, J., Lu, J.: Unipc: A unified predictor-corrector framework for fast sampling of diffusion models (2023)
1017. Zhao, X., Ananth, P., Li, L., Wang, Y.: Provable robust watermarking for ai-generated text. *CoRR abs/2306.17439* (2023). <https://doi.org/10.48550/arXiv.2306.17439>
1018. Zhao, Y., Yang, X., Wang, J., Gao, Y., Yan, C., Zhou, Y.: Bart based semantic correction for mandarin automatic speech recognition system. *arXiv preprint arXiv:2104.05507* (2021)
1019. Zheng, C., Yin, F., Zhou, H., Meng, F., Zhou, J., Chang, K.W., Huang, M., Peng, N.: On prompt-driven safeguarding for large language models (2024)
1020. Zheng, L., Chiang, W.L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E.P., Zhang, H., Gonzalez, J.E., Stoica, I.: Judging llm-as-a-judge with mt-bench and chatbot arena. *CoRR abs/2306.05685* (2023)
1021. Zheng, L., Chiang, W.L., Sheng, Y., Zhang, H.: Chatbot arena leaderboard week 8: Introducing mt-bench and vicuna-33b. <https://lmsys.org/chatbot-arena-leaderboard-week-8-introducing-mt-bench-and-vicuna-33b/> (2023)
1022. Zheng, L., Chiang, W.L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E., et al.: Judging llm-as-a-judge with mt-bench and chatbot arena. *arXiv preprint arXiv:2306.05685* (2023)
1023. Zheng, L., Chiang, W.L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E., et al.: Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems* **36** (2024)
1024. Zheng, Q., Zhang, A., Grover, A.: Online decision transformer. *Proceedings of the 39th International Conference on Machine Learning* (2022)
1025. Zheng, S., Huang, J., Chang, K.C.C.: Why does chatgpt fall short in answering questions faithfully? *arXiv preprint arXiv:2304.10513* (2023)
1026. Zheng, Z., Zhang, J., Vu, T.A., Diao, S., Tim, Y.H.W., Yeung, S.K.: Marinegpt: Unlocking secrets of “ocean” to the public (2023)
1027. Zhiheng, X., Rui, Z., Tao, G.: Safety and ethical concerns of large language models. In: *Proceedings of the 22nd Chinese National Conference on Computational Linguistics (Volume 4: Tutorial Abstracts)*, pp. 9–16 (2023)

1028. Zhong, K., Song, Z., Dhillon, I.S.: Learning non-overlapping convolutional neural networks with multiple kernels. arXiv preprint arXiv:1711.03440 (2017)
1029. Zhong, K., Song, Z., Jain, P., Bartlett, P.L., Dhillon, I.S.: Recovery guarantees for one-hidden-layer neural networks. In: Proceedings of the 34th International Conference on Machine Learning-Volume 70, pp. 4140–4149 (2017). URL <https://arxiv.org/pdf/1706.03175.pdf>
1030. Zhong, L., Wang, Z.: Can chatgpt replace stackoverflow? a study on robustness and reliability of large language model code generation (2023)
1031. Zhong, W., Cui, R., Guo, Y., Liang, Y., Lu, S., Wang, Y., Saied, A., Chen, W., Duan, N.: Agieval: A human-centric benchmark for evaluating foundation models (2023)
1032. Zhou, A., Li, B., Wang, H.: Robust prompt optimization for defending language models against jailbreaking attacks (2024)
1033. Zhou, G., Zhu, X., Song, C., Fan, Y., Zhu, H., Ma, X., Yan, Y., Jin, J., Li, H., Gai, K.: Deep interest network for click-through rate prediction. In: Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining, pp. 1059–1068 (2018)
1034. Zhou, H., Zhang, S., Peng, J., Zhang, S., Li, J., Xiong, H., Zhang, W.: Informer: Beyond efficient transformer for long sequence time-series forecasting. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 35, pp. 11106–11115 (2021)
1035. Zhou, J., Hu, M., Li, J., Zhang, X., Wu, X., King, I., Meng, H.: Rethinking machine ethics—can llms perform moral reasoning through the lens of moral theories? arXiv preprint arXiv:2308.15399 (2023)
1036. Zhou, J., Zhang, Y., Luo, Q., Parker, A.G., De Choudhury, M.: Synthetic lies: Understanding ai-generated misinformation and evaluating algorithmic and human solutions. In: Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems, pp. 1–20 (2023)
1037. Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Conditional prompt learning for vision-language models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 16816–16825 (2022)
1038. Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Learning to prompt for vision-language models. *International Journal of Computer Vision* **130**(9), 2337–2348 (2022)
1039. Zhou, T., Ma, Z., Wen, Q., Wang, X., Sun, L., Jin, R.: Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In: International Conference on Machine Learning, pp. 27268–27286. PMLR (2022)
1040. Zhou, T., Niu, P., Wang, X., Sun, L., Jin, R.: One fits all: Power general time series analysis by pretrained lm. *Advances in Neural Information Processing Systems* **36** (2023)
1041. Zhou, Y., Chu, Z., Ruan, Y., Jin, G., Huang, Y., Li, S.: ptse: A multi-model ensemble method for probabilistic time series forecasting. In: The 32nd International Joint Conference on Artificial Intelligence (2023)
1042. Zhu, C., Galjaard, J., Chen, P.Y., Chen, L.Y.: Duwak: Dual watermarks in large language models. Findings of the Association for Computational Linguistics (2024)
1043. Zhu, D., Chen, J., Shen, X., Li, X., Elhoseiny, M.: Minigpt-4: Enhancing vision-language understanding with advanced large language models. arXiv preprint arXiv:2304.10592 (2023)
1044. Zhu, K., Wang, J., Zhou, J., Wang, Z., Chen, H., Wang, Y., Yang, L., Ye, W., Gong, N.Z., Zhang, Y., et al.: Promptbench: Towards evaluating the robustness of large language models on adversarial prompts. arXiv preprint arXiv:2306.04528 (2023)
1045. Zhu, M., Tang, Y., Han, K.: Vision transformer pruning. *KDD 2021 Workshop on Model Mining* (2021)
1046. Zhu, S., Voigt, T., Ko, J., Rahimian, F.: On-device training: A first overview on existing systems (2023)
1047. Zhu, W., Liu, H., Dong, Q., Xu, J., Huang, S., Kong, L., Chen, J., Li, L.: Multilingual machine translation with large language models: Empirical results and analysis (2023)

- 1048. Zhuang, H., Zhang, Y., Liu, S.: A pilot study of query-free adversarial attack against stable diffusion. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshop (CVPRW), pp. 2384–2391 (2023)
- 1049. Zhuo, T.Y., Huang, Y., Chen, C., Xing, Z.: Red teaming chatgpt via jailbreaking: Bias, robustness, reliability and toxicity (2023)
- 1050. Ziems, C., Held, W., Shaikh, O., Chen, J., Zhang, Z., Yang, D.: Can large language models transform computational social science? (2023)
- 1051. Zong, B., Song, Q., Min, M.R., Cheng, W., Lumezanu, C., Cho, D., Chen, H.: Deep autoencoding gaussian mixture model for unsupervised anomaly detection. In: International conference on learning representations (2018)
- 1052. Zou, A., Wang, Z., Kolter, J.Z., Fredrikson, M.: Universal and transferable adversarial attacks on aligned language models. CoRR **abs/2307.15043** (2023)
- 1053. Zou, A., Wang, Z., Kolter, J.Z., Fredrikson, M.: Universal and transferable adversarial attacks on aligned language models. arXiv preprint arXiv:2307.15043 (2023)
- 1054. Zou, D., Gu, Q.: An improved analysis of training over-parameterized deep neural networks. In: H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, R. Garnett (eds.) *Advances in Neural Information Processing Systems*, vol. 32 (2019)

Index

A

Accountability, [161](#)
Adapter, [23](#)
Adversarial learning, [211](#), [212](#)
Adversarial testing, [137](#), [142](#), [167](#), [239](#), [241](#),
[249](#)
AI lifecycle, [7](#)
Alignment, [9](#), [22](#), [167](#), [185](#), [186](#)
Attention, [17](#), [26](#), [29](#), [40](#), [48](#), [93](#)

B

Backdoor, [221](#), [223](#)
Batch normalization, [16](#)
Benchmark, [59](#), [100](#), [110](#), [113](#), [151](#), [154](#), [155](#),
[162](#)
Bi-level optimization, [139](#)

C

Chain-of-thought (CoT), [10](#), [23](#), [39](#)
Concept removal, [123](#), [127](#), [249](#)
Contrastive language-image pre-training
(CLIP), [11](#), [61](#)
Contrastive search, [199](#)
Convolution, [14](#)

D

Decoder-only, [19](#)
Detection, [168](#), [169](#), [173](#), [197](#), [200](#), [209](#), [210](#),
[212](#)
Diffusion model (DM), [11](#), [84](#), [126](#), [140](#), [221](#),
[224](#), [237](#), [238](#)

E

Emerging ability, [6](#), [150](#)
Encoder-decoder, [19](#)
Encoder-only, [19](#)

F

Fairness, [158](#), [164](#)
Fine-tuning, [10](#), [77](#), [81](#), [101](#), [185](#)
First-order optimization, [77](#)
Foundation model (FM), [3](#), [77](#), [109](#)

G

Gaussian Error Linear Unit (GeLU), [15](#)
Generalization, [31](#), [41](#), [43](#), [88](#), [151](#)
Generative AI (GenAI), [4](#)
Generative pre-trained transformer (GPT), [10](#),
[19](#), [87](#), [150](#), [152](#), [155](#), [185](#)
Governance, [11](#)
Greedy decoding, [21](#)

H

Harmfulness, [190](#)

I

In-context learning (ICL), [9](#), [23](#), [37](#), [40](#), [43](#), [48](#)
Instruction tuning, [9](#), [186](#), [220](#)

J

Jailbreak, [167](#), [169](#), [173](#), [185](#)
Judge, [190](#)

L

Large language model (LLM), 4, 20, 67, 77, 88, 90, 99, 123, 136, 149, 152, 167, 185, 195, 209
 Layer normalization, 19
 Linear probing, 57, 109
 Logit, 15
 Low-rank adaptation (LoRA), 23, 81, 193

M

Machine ethics, 160, 166
 Machine unlearning, 123, 129, 134, 136, 137, 249
 Mitigation, 168, 178, 193, 231
 Model merging, 10
 Model reprogramming, 55, 88, 92
 Multi-head attention, 18

N

Neural scaling law, 5
 Next token prediction, 20

P

Parameter-efficient fine-tuning (PEFT), 10, 23, 55, 187
 Pooling, 15
 Position encoding, 16, 26
 Preference optimization, 22, 128, 152
 Prefix tuning, 23
 Privacy, 59, 67, 68, 73, 159, 165
 Prompt engineering, 9, 23, 240
 Prompt tuning, 23, 55, 71, 93, 176, 237, 240
 Pruning, 47, 129

R

Rectified linear unit (ReLU), 15
 Red-teaming, 137, 142, 155, 167, 237, 239, 241, 249

Reinforcement learning with human feedback (RLHF), 22, 152, 167, 186

Residual block, 15
 Risk management, 11
 Robustness, 158, 165

S

Safety, 11, 140, 157, 164, 167, 176, 185, 193, 223, 237, 240
 Sample complexity, 45
 Second-order optimization, 134
 Skip connection, 15
 Softmax, 15, 17, 40, 71, 196
 Speech, 99
 Supervised fine-tuning (SFT), 22, 167
 Synthetic data, 113

T

Time series, 87, 90
 Token embedding, 16, 91, 105
 Top-k sampling, 21
 Top-p sampling, 21
 Toxicity, 157
 Transformer, 16, 25, 29, 37, 48, 61, 119
 Transparency, 160
 Trust, 11, 151, 155, 162
 Truthfulness, 156, 164

V

Visual prompting, 55, 59, 60

W

Watermark, 195, 196, 198

Z

Zeroth-order optimization, 77–80, 84, 172