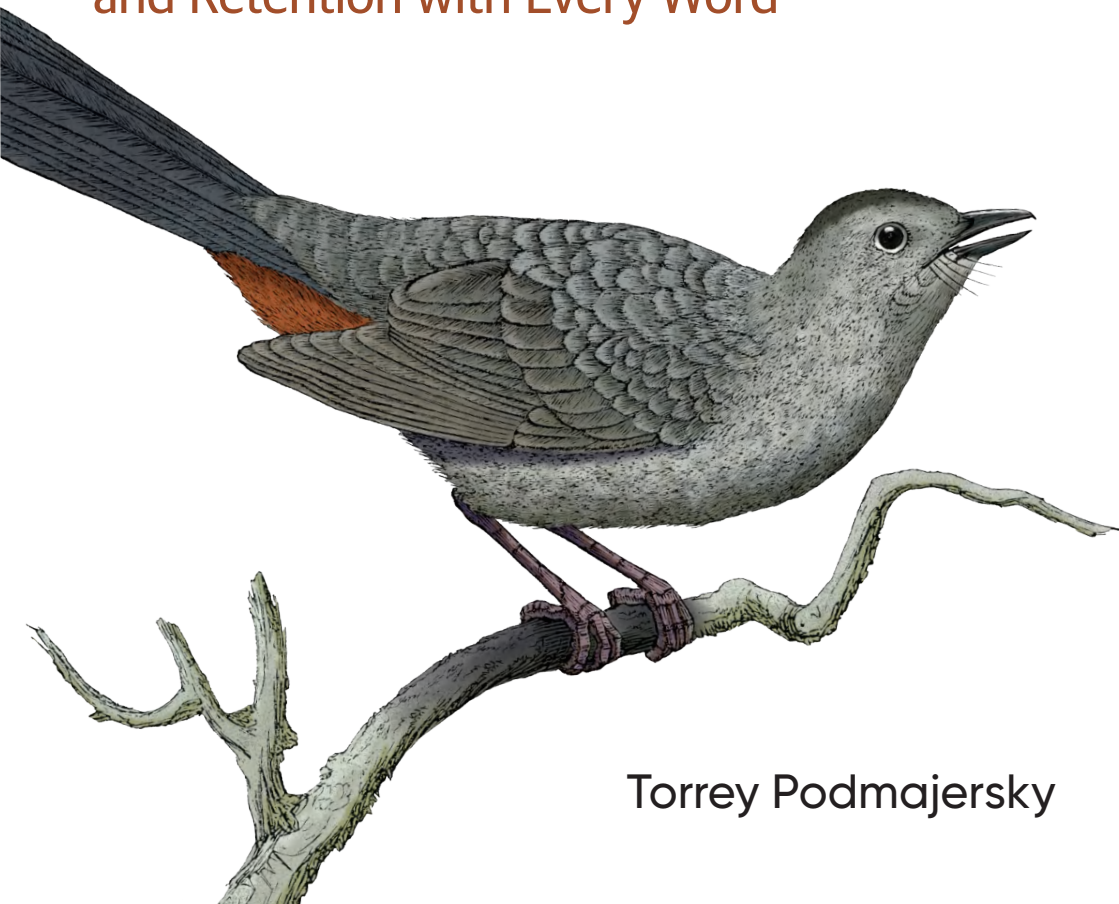


O'REILLY®

2nd Edition

Strategic Writing for UX

Drive Engagement, Conversion,
and Retention with Every Word



Torrey Podmajersky

"I found Torrey's book the one work that's actually the most realistic."

John Maeda, VP of engineering for CoreAI, Microsoft

"*Strategic Writing for UX* has long been the go-to text for understanding how to design and develop audience-centric content. This updated edition, which includes generative AI examples, is a much-needed upgrade to reflect the rapid and significant changes that have happened in content in recent years. It offers practical tools and solutions for creating UX content that can drive operational success for your organization. This will be the reference tool that you pull from your shelf again and again."

Alisa Bonsignore, author of *Sustainable Content*

Strategic Writing for UX

When you depend on users to perform specific actions—like buying tickets, playing a game, or riding public transit—well-placed words are most effective. But how do you choose the right words? And how do you know if they work? With this practical book, you'll learn how to write strategically for UX, using tools to build foundational pieces for UI text and UX voice strategy.

This second edition by Torrey Podmajersky reflects the significant advancements in the field of content design. It's an essential resource, thoroughly updated to incorporate the latest in collaborative tools, UX methods and styles, and considerations for AI-generated UX content.

- Understand the integration of UX content with the software development lifecycle
- Learn content-first design to anchor UX text in meaningful dialogue
- Apply UX text patterns effectively across varied brand voices
- Enhance product usability and conversion with strategic UI text applications
- Clearly articulate the business value of dedicated UX content efforts

Torrey Podmajersky is a UX and content strategist who has created inclusive experiences for brands like Google and Xbox and now teaches, speaks, and consults globally through her agency, Catbird Content.

DESIGN

US \$54.99 CAN \$68.99

ISBN: 978-1-098-17433-0



O'REILLY®

Praise for *Strategic Writing for UX, 2E*

“Strategic Writing for UX has long been the go-to text for understanding how to design and develop audience-centric content. This updated edition, which includes generative AI examples, is a much-needed upgrade to reflect the rapid and significant changes that have happened in content in recent years. It offers practical tools and solutions for creating UX content that can drive operational success for your organization. This will be the reference tool that you pull from your shelf again and again.”

**ALISA BONSIGNORE, AUTHOR,
SUSTAINABLE CONTENT: HOW TO MEASURE AND
MITIGATE THE CARBON FOOTPRINT OF DIGITAL DATA**

“Words are the most critical component of how we communicate product and facilitate the best possible experience. The frameworks, processes, and concepts in this book will uplevel any business that hasn’t considered UX Writing.”

MAYA JOSEPH-GOTEINER, CEO OF VELOCITY AVE.

“Torrey’s essential guide to strategic content design now incorporates the latest game-changer—AI. Every field has its foundational texts, the go-to references that establish a shared language and framework for practitioners worldwide. But few disciplines have experienced a disruption as profound as content design, with the rise of genAI. It was time for someone to update the foundation, ensuring our fundamentals remain relevant and optimized for today’s rapidly evolving landscape—thank you, Torrey!”

Yael Ben-David, Author of *The Business of UX Writing*

“I found Torrey’s book the one work that’s actually the most realistic.”

JOHN MAEDA, VP OF ENGINEERING FOR CoreAI, MICROSOFT

2nd Edition

Strategic Writing for UX

*Drive Engagement, Conversion, and
Retention with Every Word*

Torrey Podmajersky

O'REILLY®

Strategic Writing for UX, 2nd Edition

by Torrey Podmajersky

Copyright © 2025 Torrey Podmajersky. All rights reserved.

Published by O'Reilly Media, Inc.,
141 Stony Circle, Suite 195, Santa Rosa, CA 95401.

O'Reilly books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (<https://oreilly.com>). For more information, contact our corporate/institutional sales department: (800) 998-9938 or corporate@oreilly.com.

Acquisitions Editor: Simina Calin

Development Editor: Angela Rufino

Production Editor: Kristen Brown

Copyeditor: Liz Wheeler

Proofreader: Piper Content Partners

Indexer: BIM Creatives, LLC

Cover Designer: Karen Montgomery

Interior Designers: Ron Bilodeau and
Monica Kamsvaag

Illustrator: Kate Dullea

June 2019: First Edition

August 2025: Second Edition

Revision History for the Second Edition:

2025-07-22 First Release

See <https://www.oreilly.com/catalog/errata.csp?isbn=0642572041311>
for release details.

The O'Reilly logo is a registered trademark of O'Reilly Media, Inc. *Strategic Writing for UX* and related trade dress are trademarks of O'Reilly Media, Inc.

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and O'Reilly Media, Inc., was aware of a trademark claim, the designations have been printed in caps or initial caps.

Although the publisher and author have used reasonable care in preparing this book, the information it contains is distributed “as is” and without warranties of any kind. This book is not intended as legal or financial advice, and not all of the recommendations may be suitable for your situation. Professional legal and financial advisors should be consulted, as needed. Neither the publisher nor the author shall be liable for any costs, expenses, or damages resulting from use of or reliance on the information contained in this book.

978-1-098-17433-0

[LSI]

[*contents*]

	<i>Preface</i>	ix
Chapter 1	Why: Meet the Goals of People and the Organization ..	1
	Align the Goals of the People and the Organization ...	2
	Choose Content to Meet Each Goal.....	5
	Identify Purposes, Opportunities, and Constraints ...	11
	Imagine and Test Solutions.....	12
	Summary: Words Make Experiences Work.....	14
Chapter 2	Voice: They Recognize You	17
	Four Example Experiences	19
	Creating a Voice Chart.....	22
	Using the Voice Chart as a Tool	41
	Summary: Lift Every Voice	44
Chapter 3	Conversation for Content-First Design	47
	Collaborative Design.....	48
	Transforming the Conversation into an Experience...	53
	Summary: Now You're Having the Right Conversation.	56
Chapter 4	Apply UX Text Patterns	57
	Titles.....	58
	Buttons and Menus.....	64
	Descriptions	69
	Empty States.....	76
	Labels.....	81
	Controls	85
	Text Input Fields.....	89
	Transitional Text.....	93

	Confirmation Messages	98
	Notifications	101
	Errors	105
	Summary: Use Patterns as a Place to Start	111
Chapter 5	Content Design for LLM-Based Experiences	113
	LitMop: An Example of Embedded AI	114
	What LLMs Are	115
	Define the Problem	119
	Define “Good”	121
	Plan for Business and Ethical Requirements	127
	Design for Prompts and Completions	133
	Plan for LLM Measurement	143
	Adapt and Align the Results of the LLM	148
	Summary: LLMs Generate Text, but Also a Lot of Work	153
	Resources About LLMs	154
Chapter 6	Edit, Because They Didn’t Come Here to Read	155
	Editing in Four Phases	156
	Purposeful: Find Out the Job It Needs to Do	157
	Concise: Make It Short	160
	Conversational: Make It Human	162
	Clear: Make It Understandable	163
	Summary: Edit UX Text to Help People Move Forward	165
Chapter 7	Researching UX Content Effectiveness	167
	Measure UX Behaviors Directly	169
	Investigate How People Think	176
	Apply UX Content Heuristics	182
	Summary: If You Like It, Put a Value on It	207

Chapter 8	Tools and Processes	209
	Write for the Context	209
	Manage Content Review.....	216
	Document Regulated Content.....	217
	Publish the Text	219
	Track the Content Work to Be Done	222
	Prioritize What Is Urgent and What Is Important....	224
	Systematize Content Design.....	226
	Summary: The Tools Are a Means to an End.....	232
Chapter 9	A 30/60/90-Day Plan	233
	Days 0–30: What and Who.....	233
	Days 31–60: Fires and Foundations.....	240
	Days 61–90: Built for Growth.....	245
	Summary: To Fix the Words, Build Strong Foundations	247
Chapter 10	Advocate for Content Design	249
	Ground the Content in Empathy	249
	Advocate for UX Content to the Team	250
	Summary: Use UX Content to Meet Your Goals	254
	Index	255

[*Preface*]

EVERY DAY, WE MAKE thousands of decisions: buying tickets for events, playing a game, or riding public transit. Ideally, we are guided so seamlessly through the process that we barely register that we're making decisions at all. Those decisions are influenced by the language written to help you with that: the content design.

Content design is the process of creating the words in user experiences (UX): the titles, buttons, labels, instructions, descriptions, notifications, warnings, and controls that people see. It's also the setup information, first-run experience, and how-to content that gives people confidence to take the next step.

When an organization depends on individual humans performing specific behaviors like buying tickets for events, playing a game, or riding public transit, words are ubiquitous and need to be effective. Words can be seen on screens, signs, posters, and articles, as well as heard from devices and videos. The text can be minimal, but is very valuable.

But what do those words do, how do we choose them, and how do we know when they work? This book provides strategies for using content design to help meet people's goals while advancing our organizations toward converting, engaging, supporting, and reattracting those people. To meet these goals, we structure our voice throughout the content so that the brand is recognizable to its audience. We apply common UX text patterns to ease and democratize the task of writing, assess how effective the UX content is, and iterate based on our findings.

Who Should Read This Book

This book is for anyone interested in creating, promoting, evaluating or understanding the strategic importance of content created for interactive experiences.

If you need to write UX content on top of your usual job, this book is designed to help you. You might be a marketing professional, technical writer, UX designer, product owner, or software engineer. This book equips you with knowledge about what goals the UX content can accomplish, frameworks for writing it, and methods to research it.

If you are or will be a content designer (see note about job titles on page xi), or if you're a manager or leader who wants to support a content designer on your team, this book also gives you methods to demonstrate the value of content design and the impact it makes. In this book, you'll find processes and tools to do the work of writing and the work of partnering with design, business, legal, engineering, product, and other stakeholders sanely, creatively, and scalably.

How This Book Is Organized

Chapter 1 explains why UX content matters and how it integrates with the software development life cycle.

Chapter 2 provides a framework for the voice of the experience to align the UX content with the product principles.

Chapter 3 describes a process of content-first design for UX text, rooted in conversation.

Chapter 4 provides 11 patterns for UX text and demonstrates how they work in the voices of four different example experiences.

Chapter 5 introduces UX text considerations about embedding large language models (LLMs) and their dynamic generated content inside experiences.

Chapter 6 presents a four-phase process of editing UX text to be purposeful, concise, conversational, and clear.

Chapter 7 outlines three methods to measure the effect and quality of UX content: direct measurement, UX research, and heuristic analysis.

Chapter 8 recommends tools and processes for content design, including drafting text, managing content review, and tracking the work.

Chapter 9 shares my 30/60/90-day plan to ramp up and be successful as the first UX content professional in a team.

Chapter 10 concludes with advice about advocating for the content design work to be done.

EXAMPLES IN THIS BOOK

Examples throughout this book come from four fictional organizations and experiences. These are chosen to be diverse along several axes, including audience, purpose, and user context so that you can see how the same UX content advice applies in these different situations:

- The Sturgeon Club app, for members of a social club
- 'appee, a social game in which players compete by uploading images
- TAPP, an app for people who use a regional transit system
- LitMop, performance management software for employees to track their goals and responsibilities at work

TERMINOLOGY AND TITLES

For clarity, I've worked to consistently use a set of terms for the most important ideas in this book.

Content designer is the generic title I use for the team member who has the responsibility for the UX content. Other titles used in the industry include *UX content strategist*, *UX writer*, *content developer*, *content creator*, and *copywriter*, but there are many more. As with many content roles, there are nearly as many titles as there are organizations that employ us.

UX content is the output of the content designer's work that directly helps people to use the experience. *UX text* is the subset of UX content that are the words used by interfaces. Other industry names for UX text include *microcopy*, *editorial*, *UI text*, and *strings*.

Experience is the app, software, or other designed interaction the organization is building, for which the content designer is creating UX content.

Organization is the civic body, public institution, private company, or other entity that makes or commissions the experience.

Team is the group of humans a content designer collaborates with, including other UX pros, product managers, engineers, attorneys, and more.

People are the humans who use the experiences. Specific terms for people depend on the experience: people who use The Sturgeon Club are *members*, people who use 'appee are *players*, and people who use TAPP are *riders*. In LitMop, the *employees* who use the experience are different from the *customers* who buy the experience.

Why I Wrote This Book

UX content has been my professional focus since 2010. I started as a content designer for Xbox, creating experiences for the millions of people playing on the Xbox 360 console, Xbox Live, and Xbox One. Then, I worked on how people sign in to and manage their Microsoft identity and account, and was the first content designer for Microsoft Family and Microsoft Education. I was the first content designer and content strategist for the OfferUp.com marketplace, which helps millions of people buy and sell in their communities. I was the first UX content strategist for four teams at Google, working on ad measurement, ad strategy, cloud management, and customer support enablement experiences.

This book was conceived when I realized that we can't have a community or discipline of content design until we hold some basic ideas in common. Across the tech industry, we need to share expectations for what UX content can do, best practices for making the content do what it can, and methods to measure its effect. I wrote this book to share my frameworks, tools, and methods for creating effective UX content, and to share my encouragement and enthusiasm for using UX content to help people and organizations meet their goals.

Before the first edition of this book was published, we content designers didn't have a common set of frameworks, tools, or methods that addresses the unique challenges of UX content. The organizations and managers who want to hire us might know they have a "word problem," but they have a hard time figuring out who to hire, how to support us, and what impact to expect. At the same time, our tendency to do many "jobs" within our roles makes it difficult to isolate the content design portion of our work. This book will help to make it easier to identify and articulate the value of content design within our organizations.

I love making experiences that help people. For me, that includes making experiences that help people become content designers. I want more colleagues, more content designers who are developing even better methods to create great experiences. And of course, I want to help organizations improve the experience for their customers. That's why I founded Catbird Content (named after the gray catbird on the cover of this book!) to offer content design, UX strategy, and training services to organizations that want to attract and retain customers through positive user experiences.

Together, armed with common language and an understanding of core principles and best practices, we can make our digital world a better place for all.

How to Contact Us

Please address comments and questions concerning this book to the publisher:

O'Reilly Media, Inc.

141 Stony Circle, Suite 195

Santa Rosa, CA 95401

800-889-8969 (in the United States or Canada)

707-827-7019 (international or local)

707-829-0104 (fax)

support@oreilly.com

<https://oreilly.com/about/contact.html>

We have a web page for this book, where we list errata and any additional information. You can access this page at *https://oreil.ly/strategicWritingUX_2e*.

Acknowledgments

Thank you to my teams at Xbox, Windows, Microsoft Education, OfferUp, Google, and my Catbird Content clients. Everything I know about content design, I learned while working with you wonderful people. Thanks especially to those who pushed me to create better text, to find better solutions, to delight our customers and exceed the expectations of the business. I love working on challenging problems with you.

Thank you, Michelle Larez Mooney, for teaching me how to write UX. You were on my first interview loop to become a content designer, and you taught me the craft. You showed me by example how to partner effectively with engineering, product, and localization teams. Even more, you demonstrated how to engage so deeply and so effectively that the value of the work was undeniable.

Thank you, Elly Searle, for having the idea and drive to make the first content design course. You talked me into it and then went out and made it real by talking to Larry Asher at the School of Visual Concepts. I've learned so much from you about articulating what I can offer and asking for what I need. It has been a joy to teach with you and to benefit from your insights, enthusiasm, and dedication.

Thank you, Winslow Q. Hermson III, for your coaching, mentorship, and friendship. My career in UX would be neither possible nor joyful without your reflection, advice, and encouragement.

Thank you, Nathan Crowder, Jeremy Zimmerman, Dawn Vogel, Sarah Grant, and the rest of our Type 'n' Gripe. I am a writer because we wrote together every week for so many years. I wouldn't have pitched this book or finished it—twice!—without that practice.

Thank you also to my early readers and technical reviewers who helped to make this book more readable and more helpful, including Lane Goldstone, Michael Garvais, Alisa Bonsignore, Joseph True, Jennifer Bush, Toria Emery, Maya Joseph-Goteiner, Laura Klein, and Yael Ben David. Extra-special thanks to the people who provided early review and feedback on the new LLM-related content: Tom Resing, Mike Jang, and Peter Newcomb.

Thank you, Angela Rufino, Jess Haberman, Amanda Quinn, Simina Calin, and the rest of the astonishing team at O'Reilly for believing in this book, suggesting paths forward, and supporting the entire process.

And finally, thank you, Dietrich Podmajersky, my amazing partner. Your confidence that what I do matters more than housework, your support while I overcommit my time and energy, your patience for my inability to figure out when it's time to go to bed and more all add up to make the thousands of ways you made this book possible, both times. I love you.

Why: Meet the Goals of People and the Organization

If you think good design is expensive, you should look at the cost of bad design.

—RALF SPETH, FORMER CEO OF JAGUAR LAND ROVER

“WE NEED TO HIRE someone to fix the words!” I have heard this phrase from multiple people on software development and design teams I’ve worked with. In each case, the person can point to the places in the experiences we’re making where the words are “broken.” These people have recognized that fixing the words would help advance their organization or the goals of the people who use the experiences they make in some important way.

In each case I’ve seen, there is enough “fixing” to keep a person busy for years, but fixing the words will never be enough. Consider this metaphor: an experience with broken words is a house with broken walls. Fix the words as you would repair the walls.

If there’s only one broken wall, and it was built robustly, and the hole doesn’t affect the electrical, plumbing, or architectural support the building needs, we can fix it cheaply. When an experience is built with consistent terminology, voice, information architecture, and ways to find, maintain, internationalize, and update its content, all we would need to do is fix the words.

But when those things haven’t been considered, and the breaks go through electrical, plumbing, or supporting timbers, then words can’t fix the hole by themselves.

We will need a strategic approach to fix or create the underlying experience. We'll need to apply some engineering—in this case, content design—to make walls that will strengthen the whole building.

In Chapter 1, I introduce a framework for thinking about that strategic approach:

- What UX content is, and how it's different from other kinds of business content
- A framework connecting common types of content to the goals of organizations and people using experiences
- A picture of how content designers work within the organization

Align the Goals of the People and the Organization

The strategic purpose of content design is to meet two sets of goals: the goals of the organization responsible for the experience, and the goals of the people using the experience.

Let's consider the goals of our first fictional example organization: the TAPP Transit System. The TAPP Transit System is a regional transit system in a typical city anywhere in the world. TAPP, like any transit system, is under constant pressure to reduce costs and demonstrate its effectiveness. It also needs to bring in money through fares and taxes to maintain the vehicle fleet and pay its personnel.

TAPP cares about getting people to ride a bus for the first time, but that's not enough. TAPP needs to build a relationship with the region's population so that they choose to ride again, and choose to support the transit system through their political choices. The transit system needs to establish a virtuous cycle of engaging and reengaging its riders.

The cycle starts when the organization attracts people to it (Figure 1-1). Then, it needs to convert those people into riders. But because this is an app, and not just a single purchase, we need to onboard people into the experience, to set them up for success in it. Then, people can be engaged in the experience.

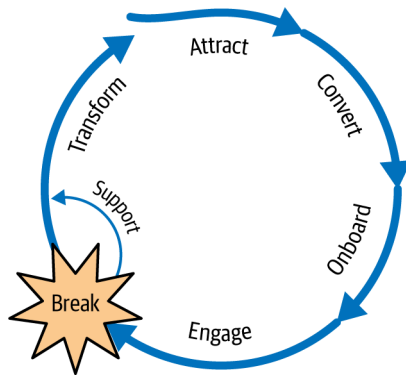


FIGURE 1-1

The organization's view of the experience virtuous cycle. Starting at the top, the organization attracts people to the experience, converts them, onboards them, and then engages them with the experience. To complete the virtuous cycle, the organization must transform engaged people into fans who attract others to the experience and who are reattracted themselves.

The virtuous part of the cycle comes next. In most cases, the organization will reap tremendous benefits if it can transform people using the experience into fans. When someone is a fan of the experience, they not only prefer to use it themselves, but they recommend it to other people, helping the organization attract new people. This transformation can happen because the experience is excellent, it's useful to them, and, like any good brand, it reflects back to the person what they want to believe about themselves.

The transformation can even happen when the experience breaks. Whatever the cause of the break (natural disaster, bad bus driver, etc.), the organization can either lose the person or support them. By supporting the people in the experience, it can retain and engage those people further. When an organization plans for potential breaks and fixes them ahead of time, it can not only continue to engage the person, but also use the break as a moment to transform a person who is merely engaged into a fan.

The local people TAPP wants to attract just want to get to work, to school, to the doctor, to the grocery store. Riding the bus might be their best option, but they need to be aware of it and trust it. They are likely unaware of the transit system's organizational goals. They probably aren't considering the variety of needs other riders might have, or

the larger goals the transit system might have. They're probably worried about all the ways their ride could go wrong: not having the right amount of money, missing a transfer, full bus, and more.

We need to understand the cycle from the point of view of the people who will use the experience, to meet them where they are (Figure 1-2). Their first task is to investigate and verify what they know about the system. They aren't expecting to be attracted into the system, and they aren't thinking about becoming part of the transit system's virtuous cycle. They just want to know their options.

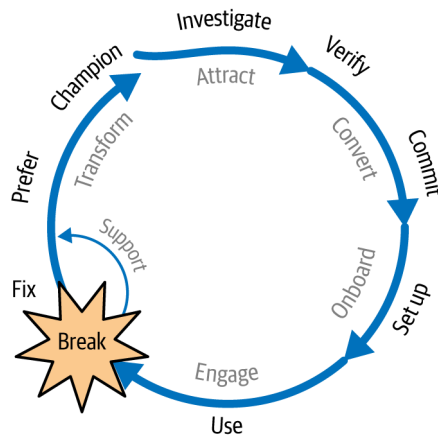


FIGURE 1-2

In the virtuous cycle, the organization and the person using the experience have different perspectives. While the organization attracts, converts, onboards, engages, supports, and transforms, the person investigates, verifies, commits to, sets up, uses, fixes, prefers, and champions the experience. By realizing this difference in perspective and focus, the organization can more effectively address what the person is there for.

Where TAPP is trying to attract the person, the person is investigating and verifying that they'll get where they want to go. Where TAPP is concerned with converting, the person is deciding or committing to the experience. Where TAPP wants to onboard and engage the person, the person expects to get on the bus, pay, ride, and arrive at their destination.

The frequent TAPP rider tends to influence their communities to ride the bus. Through their behavior, they make riding the bus seem easy. Whether they think of themselves as champions of public transit or not, these people help attract more riders into the TAPP Transit System.

Choose Content to Meet Each Goal

Throughout the virtuous cycle, content helps both the organization and the people using the experience meet their goals. What kind of content will help varies according to where the people are in the cycle.

At the beginning of the cycle, marketing content helps TAPP attract people to become riders. People interact with this marketing content to investigate and verify that the experience will be right for them. This content includes traditional marketing materials like advertisements and press releases, social media content such as videos and posts, articles written in journals, reviews and product ratings promoted on web-sites, and product pages in app stores (Figure 1-3).



FIGURE 1-3

For a person investigating whether an experience will work for them, traditional marketing content is appropriate, including ads, product pages, and more. These pieces of content meet the organizational goal of attracting people.

After a person knows about the experience, they can check whether it will work for them. To decide whether to download the TAPP app and ride the bus (or, for other experiences, to buy or download the software), a person might use endorsements, reviews, product ratings, and other types of content (Figure 1-4). All of this content helps to get people to the point of commitment.

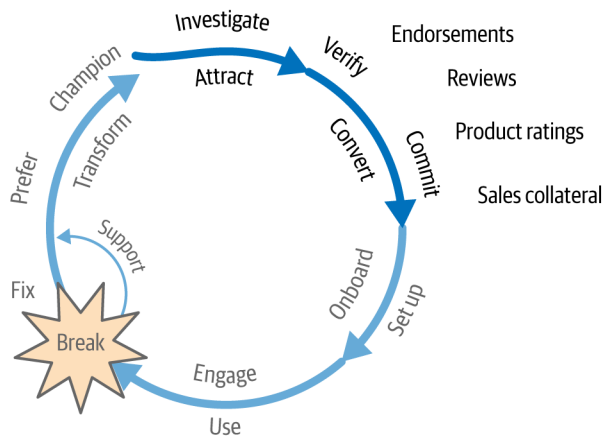


FIGURE 1-4

When a person is verifying that they should commit to trying the experience, content that helps make the sale includes external endorsements, reviews, ratings, and sales collateral.

After a person makes the commitment, marketing is over (for now.) But the app or service (the experience) still needs to be installed or set up, and the person needs to know how to take their first action (Figure 1-5). This is where UX content begins.

Consumer software like the TAPP app can require very little setup: perhaps turning on Location permission or signing in to buy bus fare. We can write UX text in the first-run experience for our TAPP app so that the first time the person uses the experience, they are able to start meeting their own goals right away.

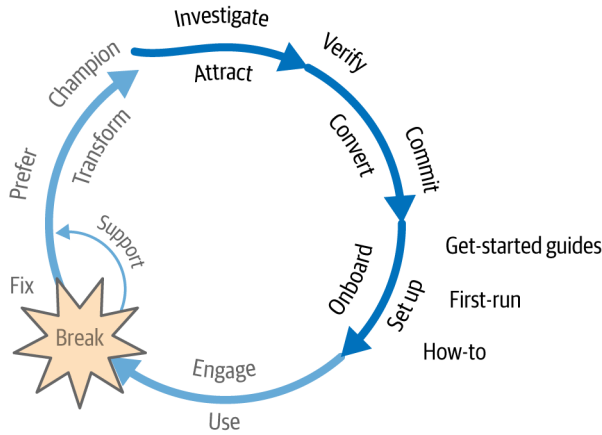


FIGURE 1-5

Onboarding helps people set up the experience. Different experiences might need different kinds of content, from simple first-run experiences to complete get-started guides and how-to information.

For software that is being used at work, there's probably more setup required. As a complication, the person who makes the decision to buy software for work is often not the employee setting up the software. At a large enough business, there are even different users: an IT pro might need to establish permissions, implement special configurations, or enter data to make the experience work for that business, before other employees use it. To meet these different goals, we need to provide different UX content for this setup crew than for the employees who will use the experience day to day.

After the experience is set up, the core UX text takes over. These words are the topic of most of this book. They are the titles, buttons, and descriptions, or voiced comments and instructions from audible experiences, that make up half or more of the interactions a person can have with an experience.

If the experience has intrinsic content, like a game, finance, or mapping app, there is special content the person is there for: the game narrative, financial information, and maps. This content might be static or dynamic, generated by LLMs or carefully designed for its purpose. To use the experience successfully, people need this content, too (Figure 1-6). For example, TAPP needs to provide static information about routes and bus fares, and dynamic timing information as buses are delayed.

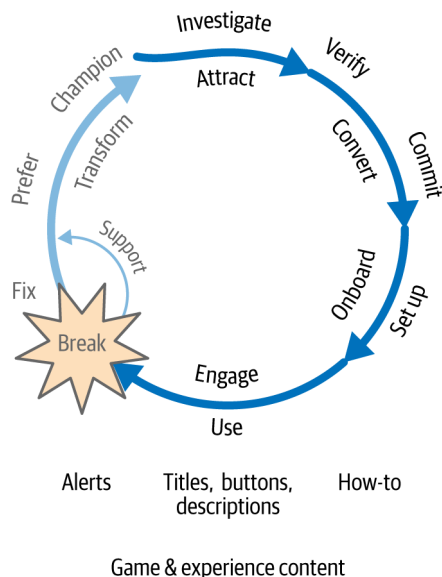


FIGURE 1-6

When people are using the experience, they interact with words in titles, buttons, descriptions, and other UX text, plus alerts and other game or consumable content.

How-to content still has a role, whether it's articles in a help center or built in to the UX context. Sometimes, people want a little confidence boost to take their next step. The job of how-to content is to give people that confidence and instruction when they want it.

Sometimes, using the experience doesn't go smoothly. Maybe the TAPP rider has forgotten to update their credit card expiration date, or maybe a bus has been unexpectedly rerouted for an emergency. The organization can use alerts and error messages to inform the person and help them meet their goals anyway (Figure 1-7). The person might seek troubleshooting content, which the organization might provide in a chatbot, a help center, on YouTube, or in scripts for support center personnel.

Supporting people through a broken experience can make those people into fans of the experience, but there are other ways to use content to help people prefer the experience. For example, we can give people badges for different kinds of engagement or allow them to get ratings in the experience. These badges and ratings give people something that they'd lose if they went to a different experience (Figure 1-8).

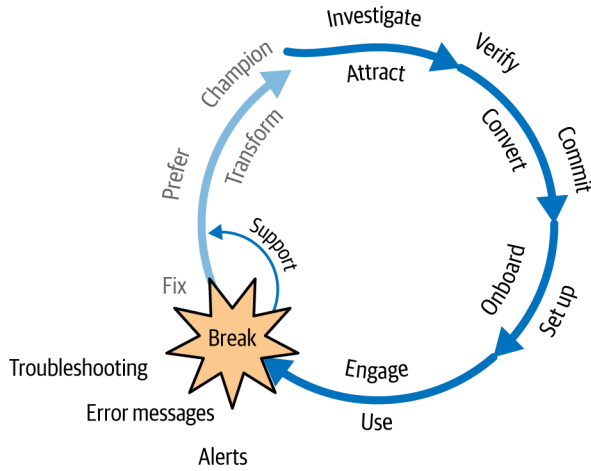


FIGURE 1-7

When there's a break in the experience, the organization can provide error messages, alerts, and troubleshooting content.

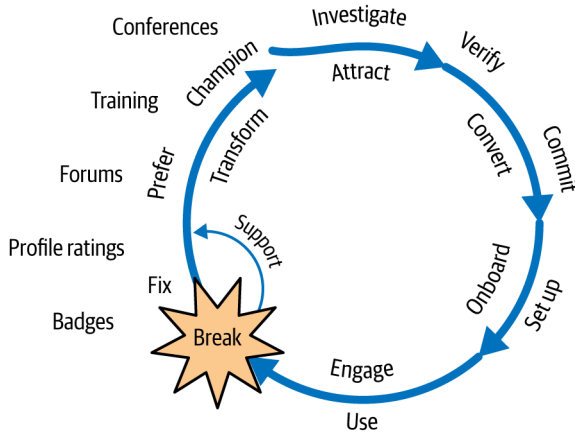


FIGURE 1-8

To give people more reasons to prefer this experience and this organization, the experience can include intrinsic value that isn't available outside of the experience and create communities around the use of the experience to help attract more people into it.

Experiences can also create communities, for example, game enthusiasts who join forums to discuss the game, or people who sell on the same online selling platform, or teachers who use a particular classroom management system. Enthusiasts of the experience join forums to share tips and tricks and to be recognized as experts.

Organizations can boost the attractiveness of their experience as well as their brand by providing forums, training, and conferences to give the fans of the experience avenues to attract new people.

Considered together, the experiences an organization makes will constitute a huge amount of content (Figure 1-9). That content is a common thread throughout the organization’s relationship with the people who buy, set up, use, and hopefully become champions for the experience.

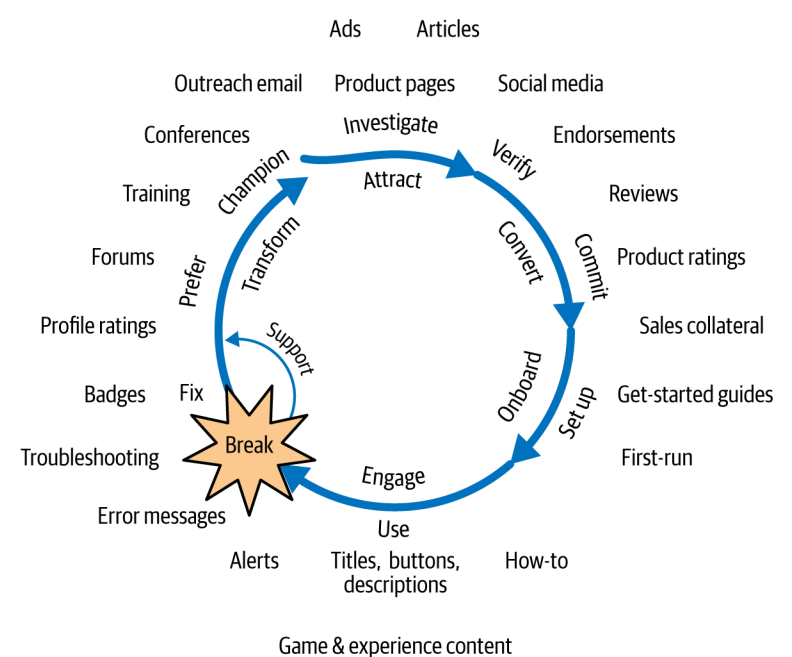


FIGURE 1-9 Examples of content that organizations use to make people aware of their experiences, bring people in, engage them, and reattract them. When the content is designed as a system, the organization benefits.

In 2018, when the first edition of this book was written, very few organizations planned their content throughout the cycle. It is more common in 2025 for consistent brand guidelines to exist because organizations

have realized something: for the person using the experience, the brand experience continues in every touchpoint. Without accessible, brand-appropriate content for onboarding, engaging, and supporting, the experience will fail to engage and transform those people into champions. Content design is how we create that content.

Identify Purposes, Opportunities, and Constraints

Writing begins where all design and engineering starts: identifying the purposes, opportunities, and constraints for the experience. Before the writing can begin, you need to identify the goals of the person who will use the experience as well as the goals of the organization making the experience.

To learn the goals for an experience, you need to collaborate with the people who understand and define those goals—the product owner, designer, marketer, researcher, engineer (not an exhaustive list)—and the people who will use the experience. From the beginning of ideation and development, you need to participate in the same meetings, discovering and defining the experience in collaboration with their team.

The primary purpose of the text is to meet the goals of the organization and people using the experience, but the text also has a role in protecting both groups. For example, the people using the experience should correctly understand how their data is used and protected. Similarly, the organization needs to have its time, money, and energy protected from liability.

From the beginning, the content designer needs to know the business strategy and constraints, including resources available for localization and the timelines to coordinate engineering and UX content with content for marketing, sales, and support. We also need to know what languages the people using the experience are fluent in, on which devices, and in what contexts. As the experience develops, we need to know technical, display, and design constraints (like maximum URL lengths and text box sizes); which text won't be updatable after the hardware is shipped; and which text can be updated from live services.

Writing for UX, just like design and coding for UX, is a design and engineering process. It is an iterative process of creation and measurement (Figure 1-10).

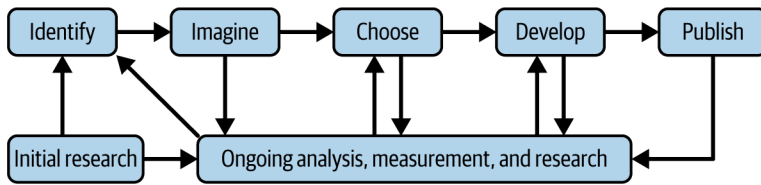


FIGURE 1-10

Writing for UX is an iterative process of creation and measurement.

To bring people into the experience, the words need to ground the conversation in what the people who will use it already understand. With the team, the content designer might conduct or collaborate on foundational, exploratory research about the context for the organization and the person who will use the experience (read more in Chapter 7).

In this initial research, the content designer can listen for sensitive topics that the experience will need to handle with care, including words that have hurtful or offensive connotations. If the experience involves money, health, privacy, government policy, or children, it's likely that complex legal or regulatory constraints apply, too. These constraints are essential to understand before designing the content for the experience.

Now that we know where we want to go, and the tools and limitations we have, we can start the most wildly creative part: imagining how to get there.

Imagine and Test Solutions

For content designers, the most creative part of the design and engineering process can be as immersive as playacting the conversation between the person and the experience, or as straightforward as adapting text that worked in the past. But whether blue-sky or mundane, the design job is to imagine and communicate at least a few distinctly different solutions. By finding several possible solutions, the team can choose the best one to move forward with.

This isn't a solo activity: even though a content designer may be responsible for marshaling the best ideas for UX content, they are not the only person with great ideas about words. Instead, the very best working groups include team members who are familiar with the technical, legal, or financial opportunities and limits, and the people who will use the experience. Those people can be experts and novices, enthusiasts

or skeptics, fans of the organization, and exclusion experts¹ who are likely to be prevented from using the experience if they aren't included from the beginning.

This working group might participate in formal design activities like design sprints, brainstorming, or the conversational design exercise that we look at in Chapter 3. They also collaborate informally, in real time and asynchronously. The content designer participates, using their special talent identifying key words and phrases. This includes helping the team discover the different terms people use, driving understanding of the ideas by clarifying words and definitions, and helping to articulate the emerging solutions in ways that the entire group understands.

The content designer may also turn to similar experiences or large language models (LLMs) for inspiration. Grounding the new content by researching what already exists in the marketplace can provide new insights the content designer can use. LLMs can mimic patterns in their training data to suggest likely acceptable phrases. However, because neither of these sources were constructed without the specific purposes of this experience in mind, they aren't a magic bullet. The content designer still needs to apply their team's understanding of the purposes for the experience, the risks, the styles, and more to imagine the right content.

After solutions are imagined, they need to be tested. Understanding what is working and isn't working about the various options is vital if the group is going to choose the best solution. From the ongoing research, content designers learn the words people already use and the phrases that resonate with them. Content designers and UX researchers can collaborate to design questions that elicit the words people would prefer to use.

In general, once the team has defined the flow and discussed information hierarchy, then UX designers (aka product designers and interaction designers) document the screens in the initial end-to-end flow. The content designer refines the design, including the language, the information architecture, and the sequence of the flow itself, in close

1 In *Mismatch: How Inclusion Shapes Design* (MIT Press, 2018), Kat Holmes defines exclusion experts as “people who experience the greatest mismatch when using your solution, or who might be the most negatively affected.”

partnership with the designers. To ensure the interactions, visual design, and text in the designs work together, we draft all of the UX text inside the designs, not as comments or separate documents. We share our best options for the content using tools that the entire team can access.

Writing UX text is iterative: you start with less-than-perfect words, then replace those words with slightly better words, and repeat until you find the best words. This is the way to make the text purposeful and protective, but also concise, conversational, and recognizable as coming from the organization's brand.

Finally, the team gets ready to launch the experience, feature, or update. Because content designers can be the single person responsible for stringing words together across all of the screens, they're often one of the very few team members with a broad yet detailed view of the whole experience. The content designer can be a big help to their support, marketing, PR, and sales partners, because they have exact and detailed knowledge of what buttons people need to press and precisely what each error message means.

Summary: Words Make Experiences Work

The value of content design is in the decisions we drive from a content perspective—the purposes, the impacts to measure, the definition of “good content” for your experiences—that add up to inform the judgment that weighs revenue against risks and costs.

In this book, I give concrete examples, tools, and advice for the content designer. But the process isn't always as clear as I've outlined in this chapter. For example, sometimes experiences are developed without clear goals in mind. Sometimes, the content designer is also the designer, or product owner, or frontend engineer, or marketer. Sometimes, the team (or individual) doesn't design several options but pursues a single vision. Most teams don't know that there's more that they can do with the words, nor do they know what to do with a content specialist.

Even if your design and engineering processes aren't ideal, I want to encourage you to consider your options for creating great, strategic UX content. If your organization or team wants to plunge forward without understanding their purpose, OK—but you can identify purposes

yourself (voice, Chapter 2). You can imagine brand-new text and the flow of an experience yourself (conversational design, Chapter 3). You can test those options with guerrilla UX research tactics or heuristics and estimate the impact the final text could make (research, Chapter 7). You can advocate for text that is conversational, concise, and purposeful (edit, Chapter 6) and write it faster by taking advantage of text patterns (text patterns, Chapter 4). You can even use the organization's experience success metrics and relate those measurements to the text (research, Chapter 7). You can streamline and systematize your own work (tools and processes, Chapter 8). And if you're just getting started with content design for your team, you can socialize the possibilities (30/60/90-day plan, Chapter 9). When you're in a leadership role that includes content design, I hope you're also championing the difference that content design can make (advocate, Chapter 10).

Voice: They Recognize You

They may forget what you said, but they will never forget how you made them feel.

—UNKNOWN, ATTRIBUTED TO MANY

PEOPLE USING AN EXPERIENCE will come away with a feeling they associate with that experience. That feeling becomes the brand of the experience and the organization that makes it. The organization wants people to remember how it makes them feel, because that feeling makes the experience recognizable, consistent, and distinct from its competitors. *Voice is the set of characteristics that allows content to create that feeling.*

As we saw in Chapter 1, the organization uses content throughout its relationship with people. When the voice is consistent throughout the virtuous cycle, brand affinity is strengthened. People can be more loyal to experiences and organizations that they recognize. Conversely, if we don't design how the content supports that feeling, the person might be left feeling anything: affection, repulsion, loyalty, disgust, or a confusion that leaves them detached.

The biggest barrier to this consistency is the potential variation among the many different team members who write the content. In large organizations, team members can be in different departments and not aware of one another. Having a common description of voice is the first step to helping diverse teams create a cohesive voice.

For example, when I worked at Microsoft in 2010, the voice for the Xbox 360 game system was “The console speaks like we’re sitting beside them, helping them play.” The “them” was well understood: console gaming enthusiasts who just wanted to play their game. How we sat beside them could be further defined: “We’re not the guy that takes the controller away and does it himself,” which could inspire disgust,

disappointment, or frustration, “but the one who will tell you exactly what to do, to make it easy for you,” to inspire feelings of camaraderie, achievement, and belonging. Because the gamer and role of the person on the sofa was so familiar to the people making the product, the definition and documentation of voice could be simple.

As Xbox started to understand its broader audience beyond the console gaming enthusiast, we adapted the voice. No matter who was playing or whether they were using the console to watch TV or listen to music, they should have a positive experience. We redefined the voice to be “Clean, casual, and keep ’em playing.” We focused the feeling on playing, achievement, and having fun.

These informal descriptions of voice are only as strong as the consistent understanding of that voice. Getting all of the team to understand that description is a major challenge because no team is monolithic. Humans will have different “feelings” for the words. Even speakers of the same language come from different regions and different backgrounds. When multiple teams need to use the same voice, those teams might be working in isolation from one another.

There are several tools emerging to help organizations create consistency in voice among their disparate teams. Some of the early ones were pre-AI, based on recognizing grammatical patterns and recommending replacements. For example, the Bethbot (<https://www.bethdunn.com/bethbot>) was created by Beth Dunn and Chris Freeley at HubSpot to create custom advice for anybody creating content for the company. Bethbot was integrated into Slack, Chrome, a chatbot, and the build system. “Any language string needed to pass the Bethbot test before it could be shipped to production”—and Bethbot gave advice about misspellings, preferred punctuation, and other detectable grammatical rules.

Some new LLM tools go a step further. These create new content or suggest edits to proffered content that align the content with a defined voice. But notice: *The voice still needs to be defined*. That’s the hard part—driving alignment among stakeholders while articulating the intangible aspects of a brand—and that’s what this chapter is about.

When we changed our definition of the Xbox voice, we put up posters in the Xbox buildings to spread the word. We created a special email address so that anybody—from development to operations—could easily reach the dedicated content design team for help. The content design team used design critiques, hallway brainstorming sessions, and peer reviews of text to stay aligned with one another.

Where there is no content design team, the process of developing and aligning the text to the voice still has to be managed across the entire organization. Even if responsibility for UX content text is centered on a single person, that person won't have enough time. UX content will sometimes need to be created without that person. When I faced that challenge at OfferUp, I created a voice chart to define the voice we wanted in a way that anybody at the company could use. You'll see several examples of voice charts in this chapter.

What you'll get from Chapter 2:

- A definition of voice and tone for written content
- The four to six aspects of text we can use to convey voice
- A method to design a product's voice to consistently convey brand principles
- Four detailed example experiences for you to leverage and modify
- Recommendations for using the product voice to reduce time to design and develop content

Four Example Experiences

When we design, our first task is to define the problem we will be solving. For content design, good questions to ask include: What kind of content will be created and for whom? Why does the user want that content, and how will they use it? These are the questions that will allow you to determine what “good” content looks like to solve this problem.

Because here's the big secret of content: how “good” content is depends on how effective it is for its purpose. Without defining the purpose, nobody can determine its effectiveness. It doesn't matter if the grammar is “perfect” (which itself has debatable meaning) if the content doesn't do its job. To design content, we have to know the job it's there to do.

To demonstrate how different “good” can look (and still follow patterns we can use), I’ve invented four different experiences. I represent each experience as it might appear on a mobile device, but The Sturgeon Club, LitMop, and TAPP could easily be web apps (Figures 2-1 and 2-2):

- **The Sturgeon Club**, an exclusive club membership app with updates about club events, reservations for facility use, dues payment, a menu, and a calendar.
- **’appee**, a casual social game with daily thematic challenges. Players compete and win prizes by uploading images according to the theme. They also rate other players’ images, make comments, and buy items imprinted with images.
- **TAPP**, a regional bus service web experience with updates per route and region. Riders can find routes, pay fares, and manage their account.
- **LitMop**, a software as a service (SaaS) product for companies to manage employee performance and allocation. Employees and managers manage their priorities and report on progress, allowing leaders to prioritize and report on human resource spending.

I designed these four example experiences to illustrate the similarities and differences in the UX content in diverse experiences:

- At work (LitMop) and at play (’appee)
- Exclusive (The Sturgeon Club and LitMop), and non-exclusive (TAPP and ’appee)
- For direct profit from users (’appee and The Sturgeon Club), for indirect profit from users (LitMop), and not for profit (TAPP)

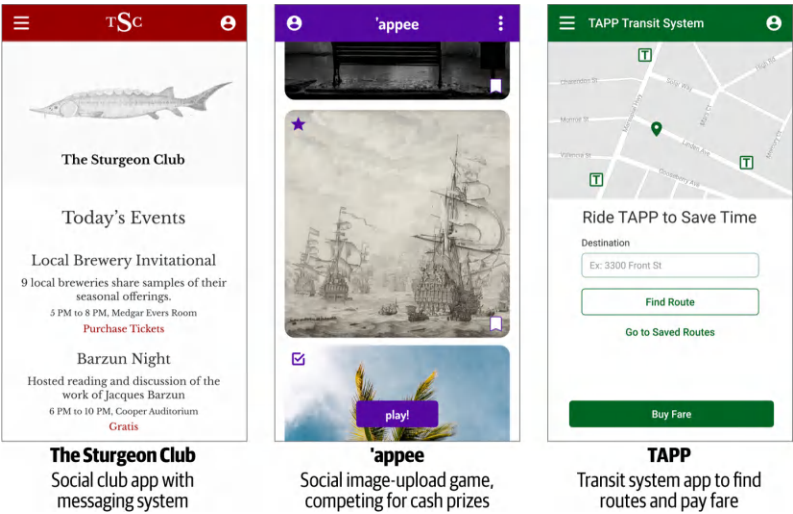


FIGURE 2-1
The Sturgeon Club, 'appee, and TAPP are three of the fictional experiences that provide examples for this book.

Dashboard

Ask HR

Hirune Matsui

Requests

Update

From Don Abene: Please update Q2 impact 3 to include target #s

Feedback

Provide for Ejiro about Q2 goals

Timeline

Jul 15

Aug 03

Aug 03

Q2 feedback requests sent

Q2 reviews complete

Q3 goals set

Goals

Q2 P1

Q2 P1

Q2 P2

Assess viability of station operations prior to acquisition...

Recommend budget parameters based on viability assessment

Preliminary cost-benefit analysis of equipment repurpose/reuse/refit

Career

Training Ladder

Setting Goals for Success

Sr. Field Researcher

Reviews

Q2 2026

Q1 2026

YE 2025

Q4 2025

Q3 2025

Q2 2025

Q1 2025

Not started

Complete: On Track

Annual Compensation Review

Complete: Exceeded

Complete: On Track

Complete: Exceeded

Complete: Onboarded

LitMop
Human Resources (HR) performance management software

FIGURE 2-2
LitMop is the fourth fictional experience created for this book.

To be clear, I'm not promoting a design vision with these apps. Their purpose is to demonstrate that even within different design choices, the same UX content principles apply—while highlighting the brand for each experience.

There is a maxim in fiction writing that each character should be recognizable from their dialog: how they speak and what they talk about. It's also a good goal that the people using an experience can recognize its voice from any piece of content. That way, when a person encounters a message or screen from the organization, they can immediately recognize it, know it's legitimate, and trust it.

In the examples from these experiences throughout this book, my goal is to make the text different enough that even if you didn't see it in context, you could tell which organization it comes from. The UX text in those examples depends on the decisions made in the voice charts created in this chapter. Let's begin!

Creating a Voice Chart

The voice chart (Table 2-1) holds a set of decision-making rules and creative guidance to make the UX content align to the needs of the organization and of the person using the experience. The voice chart has four roles: training new content designers, informing the training of LLMs, designing new text, and breaking ties among decision makers who each prefer different UX text.

When the UX content is being drafted, the voice chart will help identify what might make it better. When there are multiple good options for the UX content, the voice chart will make deciding between those good options easier. When the UX content is complete, the voice chart will help people move away from subjective judgments and use it as an objective standard (Chapter 7). When there are multiple teammates creating UX content, the voice chart helps them align UX content to the voice. I'll explain how to use it after we build it.

The voice chart shown in Table 2-1 holds each product principle (defined in the next section) in a column. Then, for each principle, each of the six aspects of voice is defined in a different row: the concepts, vocabulary, verbosity, grammar, punctuation, and capitalization.

TABLE 2-1. Blank example voice chart

	PRODUCT PRINCIPLE 1	PRODUCT PRINCIPLE 2	PRODUCT PRINCIPLE 3
Concepts			
Vocabulary			
Verbosity			
Grammar			
Punctuation			
Capitalization			

The definitions in each column relate to the product principle at the top of each column. The definitions in one column aren't expected to be the same as the definitions in another column. It's expected that even in the same row, two columns might contradict or complement each other.

This variation between columns is the difference between *voice* and *tone*. Voice is the consistent, recognizable choice of words across an entire experience. Tone is the variability in that voice from one part of the experience to another. For example, when I overhear my mother answer a phone call, I can quickly tell by her *tone* whether the phone call is from a stranger or a loved one—but I am never confused that it is my mother's *voice*. Similarly, we should be able to recognize an organization or an experience by its voice, even when the tone varies to accommodate, for example, an error message, a notification, or a moment of celebration.

By encapsulating these variations in the same voice chart, the content designer is equipped to intentionally include and vary the tone to align the UX content with the overall voice.

In the rest of this chapter, we fill in the voice chart. We begin with the product principles and then tackle the other aspects of voice: concepts, vocabulary, verbosity, grammar, punctuation, and capitalization.

PRODUCT PRINCIPLES

The foundations of the voice chart are the product principles. These principles define what the experience is trying to be to the people who use it. Then, the voice can do its job of conveying those product principles with every word.

To be clear, identifying a product or organization's principles can be, but isn't typically, the job of the content designer. If your organization has marketing or advertising support, these groups might already have defined these. There are also many skilled facilitators who can be brought into your company to facilitate. When I have facilitated the articulation of these principles, it has helped me to keep in mind that my goal (whether internal or as a consultant) is not to "own" the product principles but to help them exist so that the UX content can be aligned to them after they are ratified.

When the organization hasn't defined its principles and the role falls to you, I recommend interviewing people. In *Nicely Said* (<https://oreil.ly/LKNg8>), Nicole Fenton and Kate Kiefer Lee outline a process of interviewing people inside the organization to determine the goals of the brand, the organization, and the experience. If you are fortunate enough to have a UX research partner, this is a great place to enlist their help.

Use the results of interviews to draft the most important principles that emerge, and then ratify those with your stakeholders. As the process of refining the product principles continues, the drafts will change considerably—and that's fine. The process of articulating the product principles can become political. When I draft product principles, I expect that the first, second, and even third drafts will end up in the trash heap. The important thing is that the conversations continue, and these early drafts help the stakeholders get to the organization's goal.

For the examples in this book, I've invented three product principles for each organization. Three is not a magic number for product principles; your organization might have more or fewer.

The Sturgeon Club

The purpose of The Sturgeon Club is to provide a private, elegant venue for its membership to socialize and recreate. To bring that purpose to life, the club's executive and operations leaders have determined that the physical building, the internal spaces, and each moment the members spend with the club should be imbued with elegance, build camaraderie, and connect members to the club's traditions.

Table 2-2 shows the top row of the voice chart for The Sturgeon Club, which uses each of those three product principles as column headings: Imbued with elegance, Build camaraderie, and Connect to tradition.

TABLE 2-2. The Sturgeon Club voice chart product principles

	IMBUED WITH ELEGANCE	BUILD CAMARADERIE	CONNECT TO TRADITION
--	-------------------------	----------------------	-------------------------

'appee

The purpose of 'appee is to create an entertaining, engaging experience for its players while they generate content for the platform, engage with advertising, and buy merchandise. Instead of competing with “serious” art experiences, it is trying a strategy of playfulness, seeking to provide surprising entertainment and moments of insight.

Thus, the header of the 'appee voice chart (Table 2-3) is the product principles: Playful, Insightful, and Surprising.

TABLE 2-3. The 'appee voice chart product principles

	PLAYFUL	INSIGHTFUL	SURPRISING
--	---------	------------	------------

TAPP

The purpose of the TAPP experience is an extension of the purpose of the regional transit system itself: move people around the region, and therefore through the online experience, in a way the public finds efficient, trustworthy, and accessible. The TAPP voice chart uses those principles as the header for its voice chart columns (Table 2-4): Efficient, Trustworthy, and Accessible.

TABLE 2-4. The TAPP voice chart product principles

	EFFICIENT	TRUSTWORTHY	ACCESSIBLE
--	-----------	-------------	------------

LitMop

The purpose of LitMop is to be an HR platform that enables hiring, performance management, quarterly goal tracking, and more. Its name reflects not only a core tenet (“labor is the means of production”) but also the founder’s favorite metaphor of a “lighted mop,” in which processes should both illuminate and clean up the messy realities of business.

The primary purpose of the LitMop experience is to enable leaders to make decisions about employee performance and allocation to business priorities based on current employee data. The customers of LitMop are corporations who want to track business goals and employee performance, but the users of LitMop are the employees and managers at those corporations.

The LitMop voice chart reflects that it is intended to be used within diverse corporate cultures. LitMop works with the customer corporations, so the software needs to support any employee those corporations have: in multiple countries, with any disability, and other considerations. The LitMop voice chart reflects its founding metaphor and this context for its use in its principles (Table 2-5): Illuminating, Clear, and Inclusive.

TABLE 2-5. The LitMop voice chart product principles

	ILLUMINATING	CLEAR	INCLUSIVE
--	--------------	-------	-----------

Now that we have the principles, let’s examine the six different aspects of voice that will be listed as separate rows in the chart.

CONCEPTS

The voice chart helps us to specify, in advance, the concepts that we think will support the product principles. The concepts are the ideas or topics that the organization wants to emphasize at any open opportunity, even when they aren’t a part of the task at hand. They are the ideas that reflect the role that the organization wants the experience to have in the person’s life.

These concepts don’t imply that the experience should endlessly discuss itself and its organizational concerns. Instead, when it can, it includes these ideas. Concepts also don’t specify the words to use; they are the ideas that should land regardless of slogans or campaigns.

The Sturgeon Club

The Sturgeon Club voice, for example, specifies using the details about togetherness and belonging (Table 2-6). For example, instead of describing a formal event space as merely “Capacity of 124 people,” the UX text could mention, “Mingle with up to 124 members.”

TABLE 2-6. Concepts aligned to The Sturgeon Club product principles

	IMBUED WITH ELEGANCE	BUILD CAMARADERIE	CONNECT TO TRADITION
Concepts	Details of finish, opulence; functional and ornamental	Togetherness, belonging, and discretion	Specific connections to club members, history, fame, and power

'appee

Concepts included in 'appee are surprising information, small delights, and coincidences (Table 2-7). For example, if 'appee finds that the color blue is featured in a person's high-scoring images, it could choose between messages of "your blue images score higher" and "blue isn't sad for you."

TABLE 2-7. Concepts aligned to 'appee product principles

	PLAYFUL	INSIGHTFUL	SURPRISING
Concepts	Small delights, avoiding grand successes; frippery	Commonalities found, especially at the intersection of ideas	Unpredictable; misdirection and difficulty can be fun

TAPP

The TAPP experience adds very few new concepts to the experience. If they are included, they are specific to supporting the operating principles: a lack of waste, rides happening on time, and the inclusion of every possible rider (Table 2-8). For example, TAPP might indicate "98% on time" for a bus route at a particular stop.

TABLE 2-8. Concepts aligned to TAPP product principles

	EFFICIENT	TRUSTWORTHY	ACCESSIBLE
Concepts	Waste no resource	Every ride on time	Rides for every rider

LitMop

The LitMop experience hopes to engage employees in accurate data representation about their employment, so it focuses on providing appropriate insights from the data to each employee while making accurate data entry more likely (Table 2-9). Because the experience will be used across industries, and used by people of any age, capability, sex, gender presentation, veteran status, etc., LitMop wants to avoid even the appearance of expecting a person to have one gender presentation versus another.

TABLE 2-9. Concepts aligned to LitMop product principles

	ILLUMINATING	CLEAR	INCLUSIVE
Concepts	Insights from data	Importance of accuracy	Each employee is valued for their work

VOCABULARY

Where specific words can support or undermine a voice principle, use the vocabulary row to specify them. If there aren't specific words that help land the principle, you can omit this row.

This vocabulary row doesn't replace a robust *word list* or *terminology list*. A word list is a traditional part of a style guide to define spelling and usage choices like "canceled" versus "cancelled." A terminology list defines the words that are given meanings specific to the experience. In comparison, this vocabulary row in the voice chart specifies only the few words that are so important to the experience that they help identify its personality.

The Sturgeon Club

The Sturgeon Club vocabulary serves to reinforce the social order (Table 2-10). A member might have an *appointment* with staff such as a nutritionist or concierge. But members *meet* with one another. Generalities are to be avoided, and so is referring to someone as a "former member."

TABLE 2-10. Vocabulary aligned to The Sturgeon Club product principles

	IMBUED WITH ELEGANCE	BUILD CAMARADERIE	CONNECT TO TRADITION
Vocabulary	Avoid generalities ("very," "really," etc.)	Secure, not safe Meet with members Appointment with staff	Member Member emeritus or member (deceased), not former member

'appee

Vocabulary isn't the same kind of tool in 'appee as it is in The Sturgeon Club. In Table 2-11, Playful and Surprising don't specify any vocabulary to use or avoid. In the one principle that does reference vocabulary, it is vague but important: use plain, nonmetaphoric descriptions, such as "Your Wednesday photos are your best photos."

TABLE 2-11. Vocabulary aligned to 'appee product principles

	PLAYFUL	INSIGHTFUL	SURPRISING
Vocabulary	{Not terminology-specific}	Plain, non-metaphoric descriptions	{Not terminology-specific}

TAPP

In Table 2-12, the TAPP voice chart specifies words that could be used throughout the experience. Notably, the Accessible principle says to never use "disabled" or "invalid," but encourages the use of "available," "easy," and "ready." In practice, this means that the team avoids words that can exclude people who use wheelchairs and other assistive devices; instead, it includes them by specifying what is and isn't available, easy, or ready.

TABLE 2-12. Vocabulary aligned to TAPP product principles

	EFFICIENT	TRUSTWORTHY	ACCESSIBLE
Vocabulary	Fast, save time, save money	Regular, on time	Available, easy, ready Never use: disabled, invalid

LitMop

In the LitMop voice chart, the vocabulary row reflects the dynamic content available in the experience (Table 2-13). Because LitMop generates UX content using an LLM, the content designers are working in an ambiguous, predictive space. The illuminating column indicates that the team needs to set up an LLM to use vocabulary that is appropriate to the customer's priorities, even though the team can't know those priorities and will never see that content.

In the Clear principle, LitMop specifies maintaining consistent magnitudes for numbers. That is, they don't show \$1B beside \$1M because these are different magnitudes. Instead, they'd show \$1,000M and \$1M

to make it more likely that the numbers would be entered and understood accurately. For Inclusive, the Vocabulary row is even more prescriptive, indicating the use of singular “they” when the person’s gender isn’t available.

TABLE 2-13. Vocabulary aligned to LitMop product principles

	ILLUMINATING	CLEAR	INCLUSIVE
Vocabulary	Dynamically use customer priorities in insights	Consistent magnitudes for numbers	Singular “they” when gender presentation isn’t available

VERBOSITY

For strict usability, the words inside an experience should get out of people’s way. The UX text isn’t there to be savored or read for pleasure. But the use of too few words can be just as much of an obstacle as the use of too many; both can block a person from moving forward. Screen size and reading format makes a difference, too: people are more willing to read on a desktop computer or mobile device than on a TV screen.

The Sturgeon Club

The Sturgeon Club intentionally sets a measured pace. It is not afraid to take time to expand upon its own glory, so it will enhance descriptions with adjectives and adverbs (Table 2-14). The club also wants an air of formality, even where a more casual atmosphere is common, so it will use complete sentences (and therefore more words), even where people would usually use short phrases. However, there is a tension between setting a stately pace and wasting the members’ time: members are there to build camaraderie with one another, not with the concierge, the staff, or the experience.

TABLE 2-14. Verbosity aligned to The Sturgeon Club product principles

	IMBUED WITH ELEGANCE	BUILD CAMARADERIE	CONNECT TO TRADITION
Verbosity	Enhance responses and descriptions with adjectives/ adverbs	Be brief and begone; they aren’t here to talk to the concierge	Complete sentences even where phrases are more common

'appee

'appee demonstrates playfulness with its entry for Playful on the Verbosity row of its voice chart (Table 2-15). As a casual game, 'appee needs to introduce difficulty or challenge; without challenge, it's not much of a game. One way it can do this is by using fewer words than strictly necessary to get its point across. This cell in the 'appee voice chart is a good reminder that the voice for any experience is used like a spice when cooking: too little and the food is unappetizing; too much, and the food is inedible. If the content designer applied this part of the voice too heavily, there would be no words in the experience at all!

TABLE 2-15. Verbosity aligned to 'appee product principles

	PLAYFUL	INSIGHTFUL	SURPRISING
Verbosity	Fewer than strictly necessary	{Not verbosity-specific}	{Not verbosity-specific}

TAPP

The Verbosity row of the TAPP voice chart exhorts the team to avoid unnecessary adjectives or adverbs, except to ensure the person's success, to be accurate, and to be unambiguous (Table 2-16). As a public service, the TAPP voice aligns neatly with its utilitarian purpose.

TABLE 2-16. Verbosity aligned to TAPP product principles

	EFFICIENT	TRUSTWORTHY	ACCESSIBLE
Verbosity	No adjectives or adverbs except to ensure rider success	Enough words to have accurate information	Enough words to have unambiguous information

LitMop

For LitMop, the Verbosity row (Table 2-17) says that the length of the LLM-generated content should match the storytelling purpose it will be used for. That's pretty vague, giving the design team necessary flexibility while still implying that it should be kept short ("limit.") For Clear, LitMop works to keep content brief. For Inclusive, however, LitMop allows the use of more words when fewer words might exclude some employees.

TABLE 2-17. Verbosity aligned to LitMop product principles

	ILLUMINATING	CLEAR	INCLUSIVE
Verbosity	Limit generated content to be appropriate for storytelling purpose	Brief	Use more words when common labels are imprecise, like “person with care responsibilities” vs. “parent or guardian”

GRAMMAR

Natural language gives us a rich variety of ways to construct and convey our ideas, but all of those ways don’t work in all experiences. To maximize usability, simple grammatical structures work best for most purposes. In English, this means simple subject-predicate sentences or verb-object imperative directions such as “The bus accepts correct change and transit passes” and “Add money to your transit pass.”

However, merely maximizing usability can result in a robotic, impersonal tone. By choosing the sentence structures and other grammar that support the product principles, you have an opportunity to define the right balance of usability and personality for the experience.

The Sturgeon Club

The Sturgeon Club reinforces its culture in the Grammar row of its voice chart (Table 2-18). To imbue elegance, the experience should consider complex sentence structures. But to build camaraderie, simple grammar is preferred when discussing people. Most importantly, the club itself is spoken of in the grammar associated with formality: passive voice, past tense, and complex sentences.

TABLE 2-18. Grammar aligned to The Sturgeon Club product principles

	IMBUED WITH ELEGANCE	BUILD CAMARADERIE	CONNECT TO TRADITION
Grammar	In descriptions of experience, prefer complex to simple or compound	When discussing people, prefer simple statements	When discussing the club, prefer passive voice, past tense, complex and compound sentences

'appee

In contrast to The Sturgeon Club, 'appee prefers using the present and future tense. Even when it presents its rules, it doesn't use complete sentences (Chapter 4, Figure 4-10), as specified by the Grammar row of its voice chart (Table 2-19).

TABLE 2-19. Grammar aligned to 'appee product principles

	PLAYFUL	INSIGHTFUL	SURPRISING
Grammar	Present and future tense	{Not grammar-specific}	Phrases preferred

TAPP

TAPP continues its utilitarian style in the Grammar row of its voice chart (Table 2-20). It uses complete sentences to emphasize trustworthiness, but phrases are also acceptable, as long as they are simple.

TABLE 2-20. Grammar aligned to TAPP product principles

	EFFICIENT	TRUSTWORTHY	ACCESSIBLE
Grammar	Simple sentences or phrases	Complete sentences	Simple sentences or phrases

LitMop

For LLM-generated content within LitMop, Grammar is largely governed by the training set the model has been exposed to. This is especially important for the Illuminating principle, where it's the generated content that will provide illumination. Therefore, the Grammar line specifies that the model training data needs to include content with persuasive insights rather than any particular condition. (Table 2-21). There are no special grammatical rules for inclusive content, and simple phrases or sentences help keep the data within the experience clear.

TABLE 2-21. Grammar aligned to LitMop product principles

	ILLUMINATING	CLEAR	INCLUSIVE
Grammar	{Model trained on persuasive insight content}	Simple phrases or sentences	{Not grammar-specific}

PUNCTUATION AND CAPITALIZATION

I've heard people make the argument that punctuation and capitalization are part of the visual and typographic design of experience and are not the responsibility of the content designer. This is where most style guides do their heavy lifting: when to use commas, how to use en dashes, and more. As a starting point, the organization might choose to use an established style guide: those of the Associated Press (AP), the Modern Language Association (MLA), *The Chicago Manual of Style*, and the American Psychological Association (APA) styles, to name a few.

Even more importantly, punctuation and capitalization aren't available as design tools in some languages. For example, there are no "capital letters" when writing Chinese, and German requires capitalization for nouns. Specific punctuation is strictly required in some languages, and other languages have no punctuation as the concept is understood in English.

English compounds this problem by making its "rules" for punctuation and capitalization into mere suggestions. English can be wildly flexible. In my experience, creating a voice chart for the language the experience is developed in—for me, that's been English—has made it possible to address these differences with localization teams.

Regardless of how the style is chosen or who owns those decisions in an organization, punctuation and capitalization continue to be among the most frequent bugs I see within UX text. One of the purposes of the voice chart is to have the discussions and record the result so that future confusion can be avoided and the experience can be made consistent.

The Sturgeon Club

The Sturgeon Club voice chart details how capitalization emphasizes relationships and roles within the club (Table 2-22). It also emphasizes commas and eschews exclamation marks and tildes.

TABLE 2-22. Punctuation and capitalization aligned to The Sturgeon Club product principles

	IMBUED WITH ELEGANCE	BUILD CAMARADERIE	CONNECT TO TRADITION
Punctuation	Serial commas, colon instead of em dash, no tilde, and no exclamation mark	{Not punctuation-specific}	Sentences include terminal punctuation; titles do not.
Capitalization	Title case is used for titles, buttons, headings	Relationship roles (friend, spouse, parent) are not capitalized	Member titles, roles, committee titles, names, and roles are initial-capitalized

'appee

'appee enjoys fringe punctuation, preferring to stretch into playfulness away from tradition and formality. Instead of using capitalization to signify importance, it indicates that capitalization should be used only for emphasis (Table 2-23).

TABLE 2-23. Punctuation and capitalization aligned to 'appee product principles

	PLAYFUL	INSIGHTFUL	SURPRISING
Punctuation	Avoid periods; use emojis, exclamations, interrobangs, question marks	Tilde instead of colon, semicolon, dash, or ellipsis	{Not punctuation-specific}
Capitalization	Use capitalization only for emphasis	Use sentence case	{Not capitalization-specific}

TAPP

With its Punctuation and Capitalization rows in its voice chart (Table 2-24), TAPP focuses on clarity as the best route to efficiency, trustworthiness, and accessibility. TAPP uses commas and periods, and avoids semicolons, dashes, parenthetical remarks, and question marks. Titles and buttons are immediately recognizable as members of a hierarchy because of their capitalization.

TABLE 2-24. Punctuation and capitalization aligned to TAPP product principles

	EFFICIENT	TRUSTWORTHY	ACCESSIBLE
Punctuation	Use periods, commas. Avoid question marks. Avoid terminal punctuation for instructions.	Use periods, commas. Avoid question marks. Avoid terminal punctuation for instructions.	Avoid semicolons, dashes, parenthetical remarks.
Capitalization	Title-case titles, headings, buttons	Title-case titles, headings, buttons	Title-case titles, headings, buttons

LitMop

In LitMop’s voice chart Punctuation and Capitalization rows (Table 2-25), we again see that most illuminating content is expected to be generated automatically and that its training content is in AP style. For Clear, sentences are ended with punctuation, but labels are not. The only note for Inclusive is that the preferred capitalization of a person’s name should be maintained throughout the experience.

TABLE 2-25. Punctuation and capitalization aligned to LitMop product principles

	ILLUMINATING	CLEAR	INCLUSIVE
Punctuation	{Model trained on content in AP style}	Use terminal punctuation for sentences, not in labels.	{Not punctuation-specific}
Capitalization	Sentence-case text.	Sentence-case titles, buttons, headings, except where all caps are used for some labels.	Maintain preferred capitalization of each person’s name.

COMPLETING THE VOICE CHART

With all of the rows put together, the voice chart is a formidable tool to keep the UX content focused on meeting the goals of the organization and the people who will use the experience. Each content decision can be informed and aligned to be in the same voice, no matter who is writing that content.

When the voice is defined like this, the team can identify natural points of conflict within the voice. For example, 'appee specifies nonmetaphoric descriptions for its insights but also encourages unpredictability and the use of fewer words than are usually considered necessary. In the imagining process outlined in the next section, I demonstrate how to use these tensions within the voice chart to imagine broadly different solutions and how to choose among them.

Tables 2-26 through 2-29 present the complete voice charts for The Sturgeon Club, 'appee, TAPP, and LitMop, respectively.

TABLE 2-26. The complete voice chart for The Sturgeon Club

	IMBUED WITH ELEGANCE	BUILD CAMARADERIE	CONNECT TO TRADITION
Concepts	Details of finish, opulence; functional and ornamental	Togetherness, belonging, and discretion	Specific connections to club members, history, fame, and power
Vocabulary	Avoid generalities ("very," "really," etc.)	Secure, not safe Meet with members Appointment with staff	Member, Member emeritus, or member (deceased), not former member
Verbosity	Enhance responses and descriptions with adjectives/ adverbs	Be brief and begone; they aren't here to talk to the concierge	Complete sentences even where phrases are more common

	IMBUED WITH ELEGANCE	BUILD CAMARADERIE	CONNECT TO TRADITION
Grammar	In descriptions of experience, prefer complex to simple or compound	When discussing people, prefer simple statements	When discussing the club, prefer passive voice, past tense, complex and compound sentences
Punctuation	Serial commas, colon instead of em dash, no tilde, and no exclamation mark	{Not punctuation-specific}	Sentences include terminal punctuation; titles do not
Capitalization	Title case is used for titles, buttons, headings	Relationship roles (friend, spouse, parent) are not capitalized	Member titles, roles, committee titles, names, and roles are initial-capitalized

TABLE 2-27. The complete voice chart for 'appee

	PLAYFUL	INSIGHTFUL	SURPRISING
Concepts	Small delights, avoiding grand successes; frippery	Commonalities found especially at the intersection of ideas	Unpredictable; misdirection and difficulty can be fun
Vocabulary	{Not vocabulary-specific}	Plain, non-metaphoric descriptions	{Not vocabulary-specific}
Verbosity	Fewer than strictly necessary	{Not verbosity-specific}	{Not verbosity-specific}
Grammar	Present and future tense	{Not grammar-specific}	Phrases preferred
Punctuation	Avoid periods; use emojis, exclamations, interrobangs, question marks	Tilde instead of colon, semicolon, dash, or ellipsis	{Not punctuation-specific}
Capitalization	Use capitalization only for emphasis	Use sentence case	{Not capitalization-specific}

TABLE 2-28. The complete voice chart for TAPP

	EFFICIENT	TRUSTWORTHY	ACCESSIBLE
Concepts	Waste no resource	Every ride on time	Rides for every rider
Vocabulary	Fast, save time, save money	Regular, on time	Available, easy, ready
Verbosity	No adjectives or adverbs except to ensure rider success	Enough words to have accurate information	Enough words to have unambiguous information
Grammar	Simple sentences or phrases	Complete sentences	Simple sentences or phrases
Punctuation	Use periods, commas. Avoid question marks. Avoid terminal punctuation for instructions.	Use periods, commas. Avoid question marks. Avoid terminal punctuation for instructions.	Avoid semicolons, dashes, parenthetical remarks.
Capitalization	Title-case titles, headings, buttons	Title-case titles, headings, buttons	Title-case titles, headings, buttons

TABLE 2-29. The complete voice chart for LitMop

	ILLUMINATING	CLEAR	INCLUSIVE
Concepts	Insights from data	Importance of accuracy	Each employee is valued for their work
Vocabulary	Dynamically use customer priorities in insights	Consistent magnitudes for numbers	Singular “they” when gender presentation isn’t available
Verbosity	Limit generated content to be appropriate for storytelling purpose	Brief	Use more words when common labels are imprecise, like “person with care responsibilities” vs. “parent or guardian”
Grammar	{Model trained on persuasive insight content}	Simple phrases or sentences	{Not grammar-specific}

	ILLUMINATING	CLEAR	INCLUSIVE
Punctuation	{Model trained on content in AP style}	Use terminal punctuation for sentences, not in labels.	{Not punctuation-specific}
Capitalization	Sentence case text	Sentence case titles, buttons, headings, except where all caps are used for some labels	Maintain preferred capitalization of each person's name

To understand how voice helps make an experience recognizable, consider the sign-in screens for The Sturgeon Club, 'appee, and TAPP side by side (Figure 2-3). (You will find more examples of UX text that reflects the voice for each app, including LitMop, in Chapter 4.)

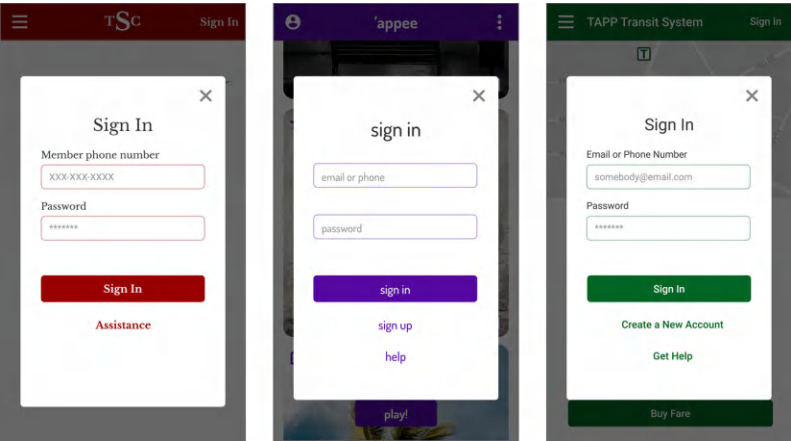


FIGURE 2-3 Sign-in screens for three of the example apps illustrate their differences in voice.

Note that the design system—the visual layout—for three of the experiences is nearly identical, but the text is a critical part of what makes them recognizably different. The Sturgeon Club specifies “Member phone number” and offers “Assistance.” ‘appee uses fewer words than usual by deciding not to offer labels for the text entry fields (a difficult choice for usability). TAPP creates the most accessible experience it can, with unambiguous, complete labels and buttons.

Using the Voice Chart as a Tool

To make the voice chart authoritative in the organization, it must be ratified by parties at the highest appropriate level in the organization. It will need sponsorship and support for the whole team to be made aware of it and to take it seriously enough to realize its value in their own work.

Ceremonies and unveilings are how organizations indicate their level of investment in an idea. To be effective and visible as a decision-making tool, the voice chart needs that investment. Plan the ceremony of a high-level sign-off.

In the meeting, walk decision makers through the voice chart, piece by piece. Provide before-and-after examples of content that can be made better by rewriting it for alignment. Demonstrate how you will use the voice chart to inform decisions and how you will measure its effect on sentiment, engagement, or other metrics relevant to your organization (see Chapter 7).

Plan a second meeting to present the voice chart to the team, and follow up by driving awareness in newsletters, email announcements, or other channels appropriate to the team's culture.

After the voice chart is adopted by the organization, it's time to use it as a tool to make decisions and make improvements. The voice chart has four roles: training new content designers, informing LLM training, designing new text, and breaking ties.

TRAINING NEW CONTENT DESIGNERS

One of the things a content designer needs to do when they join a team is to internalize the ideas, vocabulary, and grammar that the experience uses strategically. The voice chart gives them a structured reference to learn that voice in the same way they learn any other aspect of the organization.

Feedback from others is especially helpful when onboarding new team members. Using the voice chart to ground that feedback can help them to learn faster. For example, "Our voice is to use the simplest possible grammar. Is there a way to make it simpler?" Or, "Could you add more about this concept, given that it's part of our voice?"

INFORMING LLM TRAINING

Good prompt engineering for LLMs could include the desired voice chart as a set of rules. It could use those rules as a way to check existing text, though its output always needs to be checked for accuracy.

To improve an LLM’s use of voice, humans can leverage the chart to select and create content for the LLM’s underlying data set, providing good examples of content that reflects the principles intended. The chart can also be useful for humans giving feedback to the models by rejecting generated content that fails to reflect the voice.

DESIGNING NEW TEXT

The primary application for the voice chart is to design new UX text. Choose one of the product principles that applies to the moment in the experience, and draft the UX text to amplify that principle. Then, putting aside that option, repeat that drafting process with a second product principle, using those different ideas, vocabulary, and grammar.

For example, the main screen of the TAPP experience includes a map with the person’s location, a search box to find a transit route, and a main button to buy or pay bus fare. The title needs to introduce the TAPP value and promise and not distract from the main actions the person will take, either to find a route or to buy a bus fare. By using the voice chart to guide the iterations, I’ve created three versions of the title, one aligned to each of the three principles (Figure 2-4).

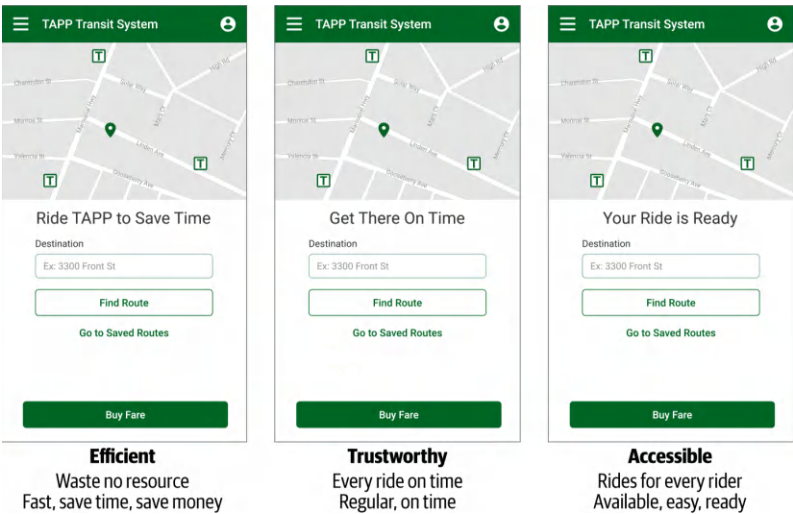


FIGURE 2-4

The main screen of the TAPP Transit System experience shows different possible headlines to align with each of the TAPP principles.

By creating versions of the content that align to different product principles, each of which are a part of the brand, we're exercising the content. The UX text becomes capable of lifting more weight and is more capable of meeting its purpose when we have a clear articulation of that purpose and how we intend to meet it. That's what the voice chart is for.

For any set of text, this iterative process will give you a broader range of options to choose from. The more varied the options you can share with your team, with all of them at high quality, the more you will shift the conversation from "We need to fix the words" to "Let's find and test the best options." Then, it's time to decide which of the good options to use.

DECISION MAKING AND TIE BREAKING

When you have created several good options for the UX text, you have multiple good options you can test against one another to determine the difference in their effectiveness (read more in Chapter 7 about research).

When testing isn't possible, practical, or desirable, the team can usually nominate one of the options as their favorite. When there are disagreements about which option to use, especially for experiences that don't have usage data available, breaking a tie will come down to how your organization makes decisions. There are three common methods that I've seen in teams and organizations: consensus decisions, autonomous decisions, and hierarchical decisions:

Consensus decisions

When an organization has a preference for driving consensus, make the case for the best option. Frame your argument by illustrating the problems to be solved by the design, including both the immediate needs and the broader organizational objectives. Use the voice chart to remind the group about the organization's need to build the brand's relationship with the people who will use the experience.

Autonomous decisions

When an organization prefers independent, individually responsible work, the choice might be all up to you! In addition to seeking feedback from others, use the voice chart as your own personal checklist: does the text include the right ideas? Is it phrased according to the predefined grammar? If you have two great options for voice and usability (Chapter 7 includes heuristics for usability), either of them will work. You can literally flip a coin.

Hierarchical or autocratic decisions

In many organizations, people higher in the hierarchy are the designated decision makers, regardless of who holds specific knowledge or expertise. The decision makers want to make the best possible choice, both for the organization and the team, so they will seek information from their experts and networks. In the ideal situation, the decision maker trusts their experts; otherwise, their decision-making powers can't scale. By consulting with their experts and networks about the benefits and risks of the options, the decision maker and the organization can have more confidence in their decision.

If the option preferred by the decision maker isn't aligned to the voice, the voice chart itself can serve as a tiebreaker. Because you did the work to have the voice chart ratified, it holds the same authority as the highest-level person who signed off on it. For example, a team that disagrees with the voice needs to have a rationale that will convince the CEO, if the voice was signed off by the CEO. If you decide and agree to fundamental terms in moments of calm, it's easier to fall back on them in moments of chaos.

Summary: Lift Every Voice

The voice of an experience is made up of many choices in the text. It begins with the ideas we choose to include or exclude, even if those words don't have a detectable difference on the "doing" at hand. It continues with the words we choose, how many we use, how we organize them, and how we use punctuation and capitalization.

When we create the voice of an experience with intention, we can use word choices as a power tool to align every word to the goals of the organization and the customer. But it's not a one-person tool: creating the voice chart is work that will take time and investment from a broad set of stakeholders.

Even if a content designer is certain that they could create the voice chart in isolation, they should resist the temptation. The minimum team to establish the voice will include representatives from marketing, research, product, leadership, support, and design. The experience will reflect the team who makes it, so we can create greater and more scalable future success by shepherding the team through the process of defining how the product principles affect the voice. To get the

experience speaking in the new voice, they will need to consider it, commit to it, and practice it. Together, the team can use the voice chart to create the feelings that people are seeking and better create the success that the organization needs.

Creating those feelings with UX content starts by writing the words that people will experience. In the next four chapters, we dive into practical techniques for writing, generating, editing, and measuring those words.

Conversation for Content-First Design

The role of the designer is that of a good, thoughtful host anticipating the needs of his guests.

—CHARLES EAMES, AMERICAN DESIGNER

WRITING FROM SCRATCH IS DAUNTING when the page is blank, the sky is blue, and the task is described only as “make something entertaining.” But that’s not what content designers are for. Our words aren’t there to be read, savored, and appreciated, but to pass unremembered while they help get somebody to the thing they want. When we approach content design, we know where we start: the goals of the organization and of the people who will use the experience, and the work we’ve already done to determine the voice.

In this chapter, I share an exercise based on the primary way humans interact with others: *the conversation*. It’s a method of designing an experience that starts before the diagrams or screens. (For working on existing UX text, try Chapter 6.) At the beginning of a new design process, we can lead with content when we start with goals (everything starts with goals!), create the conversation, and create wireframe designs from there.

Conversation is somehow in our genetic makeup. Humans take turns speaking and responding in ways that cross languages, continents, and cultures.¹ Conversation is a lot older than responding to pixels on screens and sounds from speakers, and it still governs how we respond to those pixels and sounds.

Throughout this book, when I write that UX text should be *conversational*, I am not specifying a voice or tone, like “casual conversation” or “folksy.” I mean that it is recognizable to humans as a word-based interaction they are having. It’s a pattern of interaction that respects the norms of conversation. When a person is interacting with the experience, they are in conversation with it.

What you’ll get from Chapter 3:

- Why *conversational* refers to something deeply human, beyond a tone or voice
- A method for content-first design based in conversation
- How to represent a conversation in a user interface (that isn’t a chatbot!)

Collaborative Design

In this exercise, you’ll work through an experience as if it were a conversation between two people. You’ll need to have an idea of the places people will start from when they use the experience and what they want. You’ll also need to know why the organization wants this interaction to happen and the relevant constraints that both parties are subject to. In its best circumstance, you’ll have one or more people to create these designs with, which is how I’ll describe the process.

To prepare for conversational design, prepare an online workspace like Miro or FigJam as described in the paragraphs that follow, or collect some sticky notes, a whiteboard, and markers to work in physical space. Gather your partner or small group; none of the group is required to be a dedicated content designer, but all of them should have some familiarity with the organization and the people who will use the experience.

1 For more on this, I recommend *How We Talk: The Inner Workings* by N. J. Enfield (<https://nickenfield.org/books/how-we-talk>) and also *Conversational Design* by Erika Hall (<https://www.mulebooks.com/conversational-design>).

It's not only productive but fun to assemble a group for this exercise: from 2 to 10 people, including representatives of the core stakeholders for the organization and the people who will use the experience. For your organization, that might be team members from design, research, product, business, and engineering. For the people who will use the experience, variety is key: you want a sample that can represent the breadth of the humans that your experience will try to serve. They can be people new to the experience, people who already use it, people who use similar experiences easily, and people who are excluded from similar experiences.

Note what you don't need: you don't need designs. You want to avoid thinking about the design elements that you could use to represent the conversation. All of that will come later.

Begin by identifying where the person is coming from and what outcome the person wants. For business- or security-driven features, it can be challenging (and important!) to identify these. Determining the starting place and outcome will help you capture the person's motivations to engage in the experience at all. Many experiences will have different kinds of people, with different priorities and needs, coming to do the same task—but even starting with a single person (or persona) can get you started.

In this example, we examine how to renew a bus pass with TAPP, so we put these starting and ending places on a wall or whiteboard that the entire small group can use, at either end of a long arrow (Figure 3-1). The rest of the conversation will go in the middle.



FIGURE 3-1

At the beginning of a conversational design exercise, start with the person's intention at the beginning of a long arrow and the result they want at the other end of the arrow.

Next, with input or participation from your stakeholders, make lists to answer two questions:

- Why is the person in this experience; why are they doing this thing?
- Why does this organization provide this experience?

Record the answers to these questions as the goals of the person and the organization, so that they can be referred to and updated during the exercise (Figure 3-2).

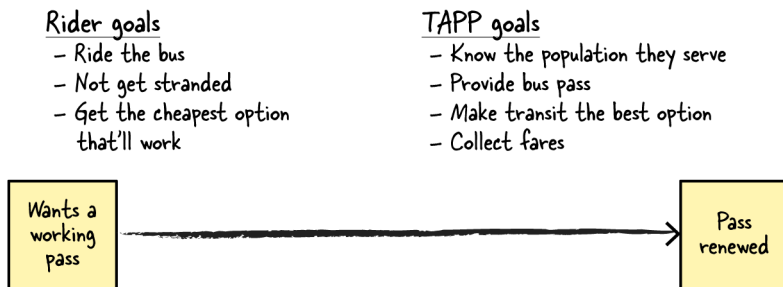


FIGURE 3-2

To inform the conversation exercise, add two lists to the long arrow: the goals of people using the experience and the goals of the organization.

Now that the basics are in place, it's time to role-play. People in your small group will act out two roles: one group member will be the person using the experience, and the second group member will embody the experience itself.

The job of the group member role-playing the experience is to help the other person achieve the outcome they're looking for while meeting the organization's goals. The experience is the thoughtful host, anticipating the needs of the person.

The job of the person using the experience is to make clear what they want and to keep pushing toward that outcome in a way that will meet their goals.

Together, both people will improvise a conversation and then iterate it until the conversation meets all of their goals. This is similar to improvisation techniques used in theater, but you don't need to perform for an audience. You'll only save the best parts to create the experience.

It's also possible to use ChatGPT or similar LLM-based technology to pretend to be one party or the other. For example, you can use the prompt "I'd like you to pretend to be an app that lets me renew a bus pass. Can you do that?" Or conversely, "I'd like you to pretend to be a person that needs to renew a bus pass, and I'll be the app." As of May 2025, these LLMs do a creditable job approximating both an app and a person using such an app (Figure 3-3), though they don't include the refinements that bring the experience to life. They also will only return results similar to what they've seen before, making them not a good partner for innovation. We'll get to those refinements in this chapter.

ChatGPT 3.5 ▾

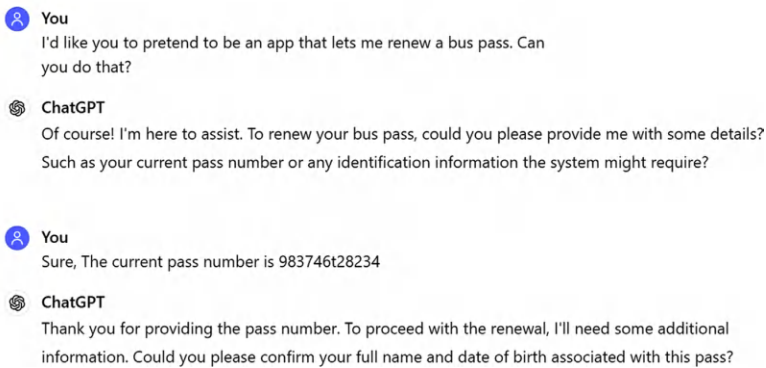


FIGURE 3-3

ChatGPT 3.5 pretends to be an app that will renew a bus pass. Screen captured May 27, 2024.

If you're working in person, instead of online or with a chatbot, the best way to role-play is to get up and move. By letting body language do some of the talking, the team will uncover nuances of the conversation that will make a better user experience—and more fun for the team! For example, if the experience involves a purchase, set up the physical scene for success by having the "organization" stand behind a counter, like a faux cashier. If you're working on the first-run or onboarding experience, start on opposite sides of a closed door, and let the person knock on the door to enter.

After the actors are physically in place, the person playing the experience starts the conversation. "What can I help you with?" is a pretty good place to begin (even if you decide to change that start later).

If you're running an activity with multiple people, each person should have an opportunity to be the experience and the person who uses the experience. A best practice is to vary the kinds of people who are represented, including different needs and different capabilities. Even LLM-based chatbots can perform both roles, though we need to remember that they are limited by the biases and age of their data sets.

Unless you have a team that has practiced improvisation, this process might feel awkward the first few times. *Stick with it.* I recommend getting through the conversation at least twice before pausing the first time. If any role-player becomes stuck, they should refer to the list of goals and outcomes. When all of the goals and outcomes are met for both parties, they can end the experience.

On the opposite end of the spectrum, if you use an LLM for this exercise, it is unlikely to express awkwardness. Instead, it may leave out information or topics that are necessary to cover that real situations might actually need.

Regardless of who is acting out the conversation, record the topics on the arrow diagram you've already started (Figure 3-4). Write the topics in the order in which they happened between the two endpoints. When people come up with great phrases to explain a concept or ask a question, write those down.

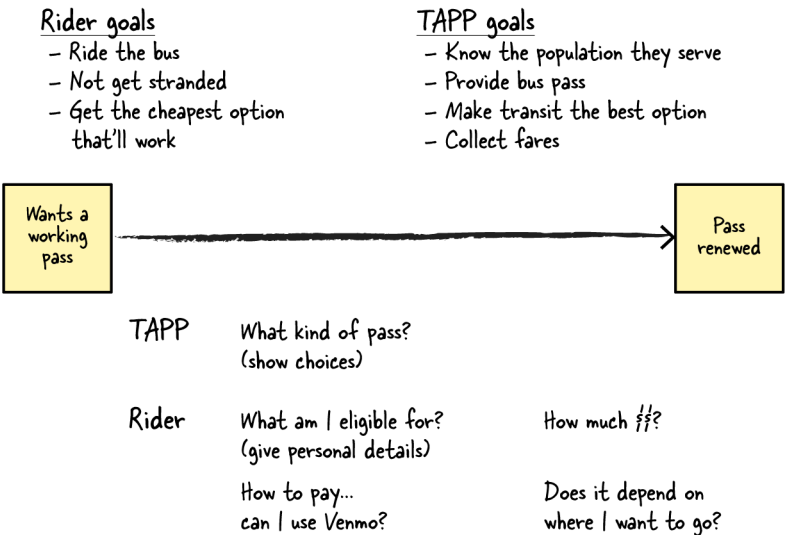


FIGURE 3-4

Rough notes, including the phrases and sequence of the conversation, are drawn where they appear on the arrow between the two endpoints.

As a team, you can consider and role-play the following: what would happen if you asked the questions in a different order? What would happen if you asked them in a different way? What if the person were a child or someone who needed something more or less complex? This exercise is your laboratory to test those ideas.

By intentionally choosing the sequence of the topics, you can make the experience more effective and even more enjoyable for the people who will use it. It also helps to raise terminology to the surface and introduce what terms might need to be defined, rethought, or specifically introduced.

At this phase, the conversation might be messy. Thinking itself is messy, and this exercise provides a structure to make sense of it. These diagrams give you the starting draft for the design work that will follow.

Transforming the Conversation into an Experience

After you have explored the sequence of topics in the conversation, varying the language you use, it's time to make it real. Take pictures or save the messy version, in case you need to refer to it later. But also distill the messy version into a clean version that you all agree on, like the row of notes on the arrow in Figure 3-5.

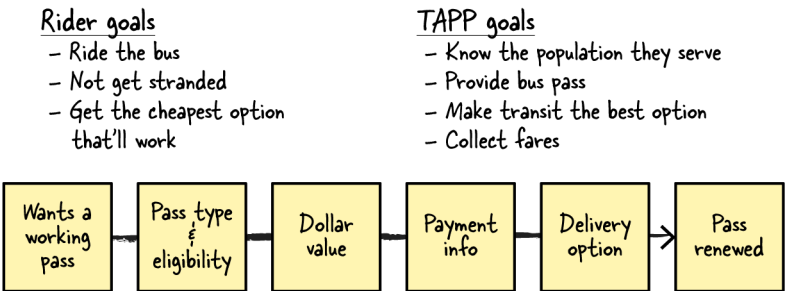


FIGURE 3-5

The rough notes of the conversation exercise are distilled into a sequence of notes, the framework of a “user journey” that meets the goals of the people and the organization.

Next, it's time to see what the words look like when written down. Some spoken words, although common and conversational, aren't easy to read. One way to write it is as a text message conversation, using side-by-side text balloons (Figure 3-6).

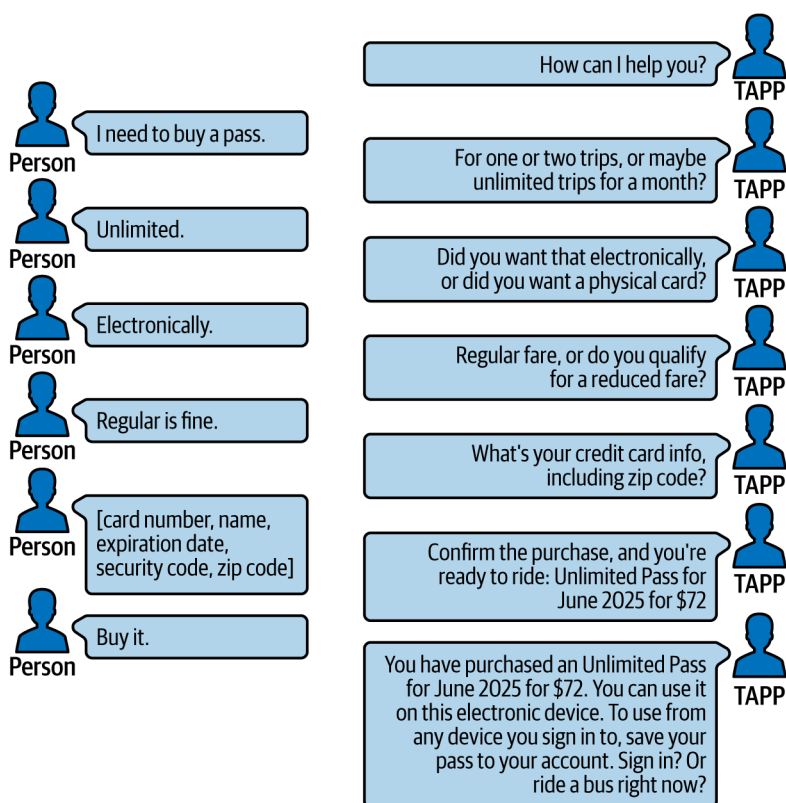


FIGURE 3-6

Here is a refinement of the conversation between TAPP and a person who wants to buy a bus pass using side-by-side text message bubbles.

At the end of a conversational design exercise, you have a design for the overall conversation. The phrases said by the experience will become the titles, labels, and descriptions. The phrases said by the person will become the buttons and options they choose within the experience. You know when key terminology needs to be introduced, and you have draft text that you can use to get started. You're unlikely to have explored all possible paths, but you have a good start!

This is enough to sketch a visual experience, to start mapping a voice interface experience, or to create a physical, in-person experience.

An initial sketch of buying a TAPP pass, using content-first design, might look like Figure 3-7.

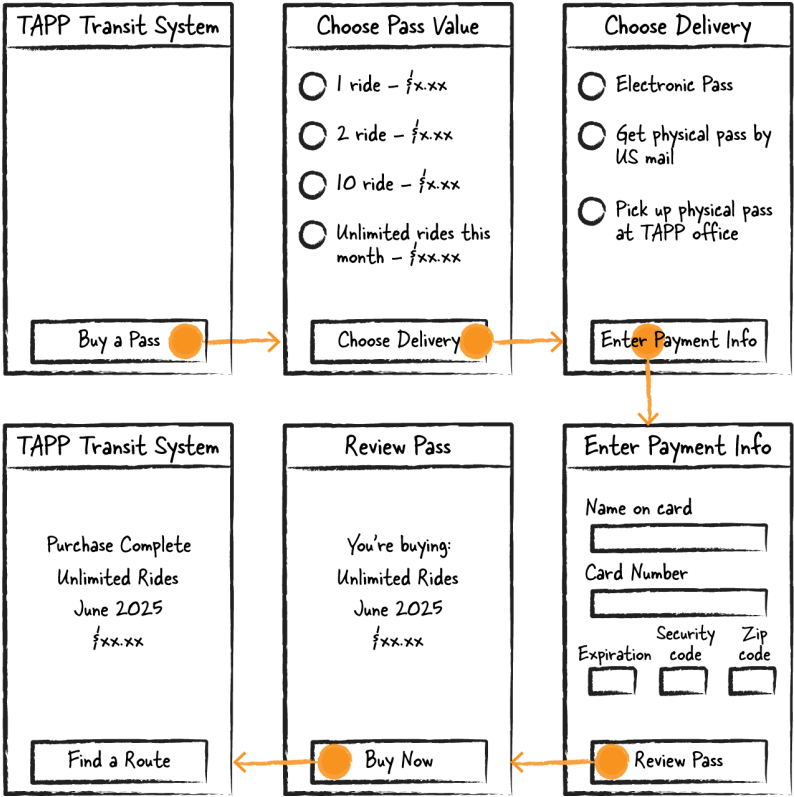


FIGURE 3-7

Here are the results of the conversational design exercise: an initial sketch of buying a bus pass in TAPP, representing all six steps.

In the exercise, the team might identify entry points to additional experiences (in this TAPP example, to sign in or ride the bus). They'll also uncover edge cases and error situations like expired credit cards. Common error conditions can be used to pressure-test the designs, making sure that the experience is robust enough to handle them. The new entry points, edge cases, and error conditions can be documented with the design so that the experience feels like a cohesive whole.

Summary: Now You're Having the Right Conversation

The UX text and design aren't optimized, but you have accomplished the most difficult, most essential work: the conversation. The team knows the experience will be aligned to the goals for both the organization and the people who will use it. The team can be confident that the right conversation will happen.

This example may seem trivial; why work this hard for such a well-understood interaction? I used an easy interaction here so that the exercise could be understood. But let's imagine a more complex scenario: enabling a systems administrator to meet compliance standards by establishing a new cloud storage solution with usage monitoring that generates automatic security reports. Even with multiple overlapping goals, inside complicated systems, this conversation design can lead to high-usability, brand-aware designs.

But the work is not complete after the wireframes are created! Now that the conversation exists, the team can refine the experience. This single path can be widened and forked according to the different needs and circumstances of the people who will use the experience. The content designer can go straight to the iterative editing process or get a jump start with those edits by applying UX text patterns.

Apply UX Text Patterns

Pay attention to the intricate patterns of your existence that you take for granted.

—DOUG DILLON, WRITER

A DESIGN PATTERN IS a reusable, common solution to a design problem. In this chapter, you'll see UX text patterns for 11 common interactions. These are arranged like a reference book, rather than a "read straight through" book. They are intended for you to use while you start drafting or editing UX text, or for you to embed within a design system for reference whenever you use a particular component.

The UX text patterns in this chapter are supported by the proprietary research I conducted and consumed at Microsoft, Google, and OfferUp while creating experiences for people to use at work and at play. All of that research was conducted in English and doesn't apply across all languages! I am hopeful that more research will be conducted and shared for other languages and cultures, but I don't have that information.

The goal of this set of patterns is to establish an easy, recognizable starting place to write consistently high-quality text. They are a tool to quickly and scalably write new UX text based on text patterns that have been successful in the past.

Like other good design patterns, these patterns don't prescribe the words to use. They also shouldn't create the impression that the pattern is necessary to solve any particular problem; sometimes, UX text isn't the appropriate solution at all.

The UX text patterns included here are a basic set that almost every experience uses:

- Titles
- Buttons and menus
- Descriptions
- Empty states
- Labels
- Controls
- Text input fields
- Transitional text
- Confirmation messages
- Notifications
- Errors

For each of the UX text patterns in this chapter, I provide three pieces of critical information: the pattern's purpose and definition and tips for its use. I also provide examples of each pattern using the book's example experiences: The Sturgeon Club, 'appee, TAPP, and LitMop. This way, you can see a variety of text patterns in different voices.

To help keep the examples clear, three of the experiences are shown as mobile apps. The Sturgeon Club screens are on the left side, 'appee screens in the middle, and TAPP screens on the right side of their respective figures. LitMop is shown as if it were on a wider computer or tablet screen. Regardless of format, the same UX text patterns can be applied to mobile, desktop, and TV screen experiences.

Let's begin with the first piece of content encountered in most experiences: titles.

Titles

A title indicates the top levels of hierarchy in the information architecture. Titles are frequently the first and only text a person reads in an experience. That means that for the person to be successful, the title needs to provide context and sometimes direct the user to take an action.

In the title, we should show the specific words that people expect to see. This way, we can increase their trust and sense of certainty that they're in the right place or taking the action that they intended to take.

To ensure the person recognizes the words in the title, even in a quick scan or listen, it's critical to keep the title brief. In English, I expect people to notice keywords best when the overall length and depth (number of lines when the text wraps) is very short.

The difference between setting context and specifying an action is a key design decision: Is it more important to reinforce where the person is in the experience or to direct them to an action they could take? In the first edition of this book, I identified three types of context-setting: using the brand name, using the content name, or naming an ambiguous task. However, now I think it's more effective to group them together as different options for context-setting.

CONTEXT-SETTING TITLE

At moments in the experience that define the brand, the experience itself is the context to set. Use the name of the experience to set that context, as a brand-name title.

For example, the main screen of The Sturgeon Club uses the club monogram and name as the brand-name title (Figure 4-1). This is the screen each member of The Sturgeon Club experience would encounter most often, and it should convey the recognizable brand.

Many apps use the very top of the screen for a title relevant to the screen, but not The Sturgeon Club. Part of what the club provides for its members is the feeling of belonging to the club itself, so The Sturgeon Club has committed to having a persistent monogram, part of its brand, acting as a title for every screen.

On the body of the screen, The Sturgeon Club reinforces the information hierarchy with a second brand-name title. By including both titles on the main page, the experience subtly sets the expectation that each screen will set its own context within the page.

Another kind of context to be set is to promote the content within the experience. When a screen is based on content, such as a blog post, social media post, or image, the screen might use a title based on that content: the content name title. These content name titles can be specified by the person who published the content, like the publisher of a blog, or they can be generated from the content itself, like the titles generated for single posts on social media.

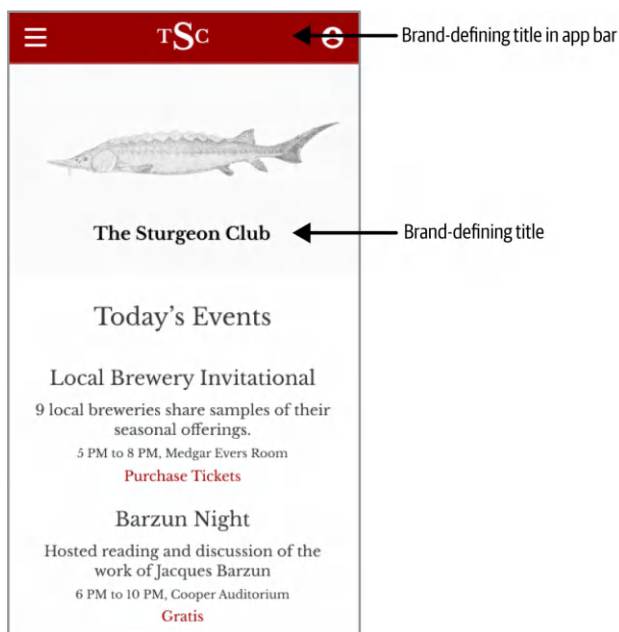


FIGURE 4-1

Brand-name titles appear on the main screen of The Sturgeon Club experience.

For example, 'appee uses the name of the challenge as the title for every image submitted for that challenge. For the 'appee screen that shows the grand-prize winner for the Bluster challenge, the appropriate title is "Bluster" (Figure 4-2).

In contrast with a title that provides the context of a single piece of content, there are some experiences that display lots of information together. On these screens, like a dashboard view of a person's account, it's helpful to use a title that covers the entire set of ambiguous tasks.

To set the context for the ambiguous (unknown to us) reasons the person might visit that screen, we can use a noun or noun phrase that names the person's context or a verb phrase that indicates the relevant category of information that the screen provides, or the variety of actions that they can take. The title can reassure the person that they are in the right place to accomplish their goal, even though the experience doesn't know which goal the person has in mind.

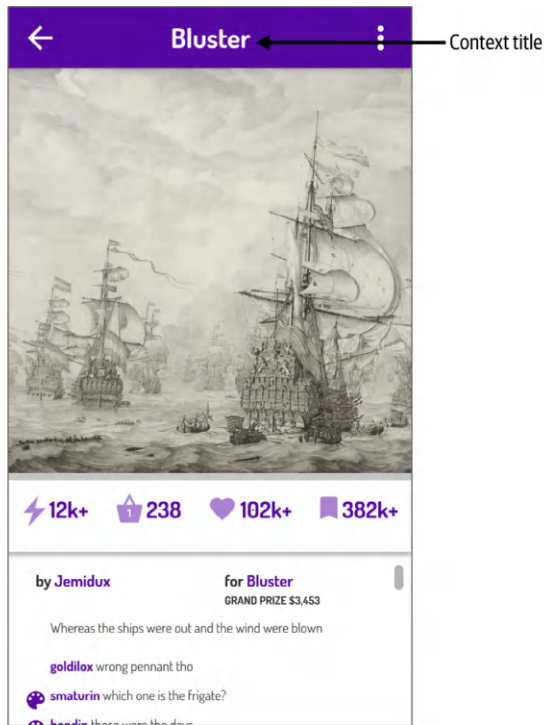


FIGURE 4-2

'appee uses the name of the challenge as the title for every image submitted for that challenge. In this figure, the winner for the Bluster challenge has the content name title “Bluster” in the app bar.

For example, there are many reasons a person might come to the LitMop dashboard (Figure 4-3). *Dashboard* is a commonly used word for a place to see the status of many different items, borrowing the term from the dashboard of a car, where one can see speed, fuel level, and alerts. It doesn't give an indication of what someone can do there, because the dashboard itself isn't an action space. Instead, it tells you what actions you might want to take elsewhere or reassures you that there are no actions to take.

These context-setting titles serve as wayfinding, helping people establish where they are and hinting at what they can do there. In contrast, when the purpose of that screen in an experience is to take a particular action, we want to help them take that action with the single-task title.

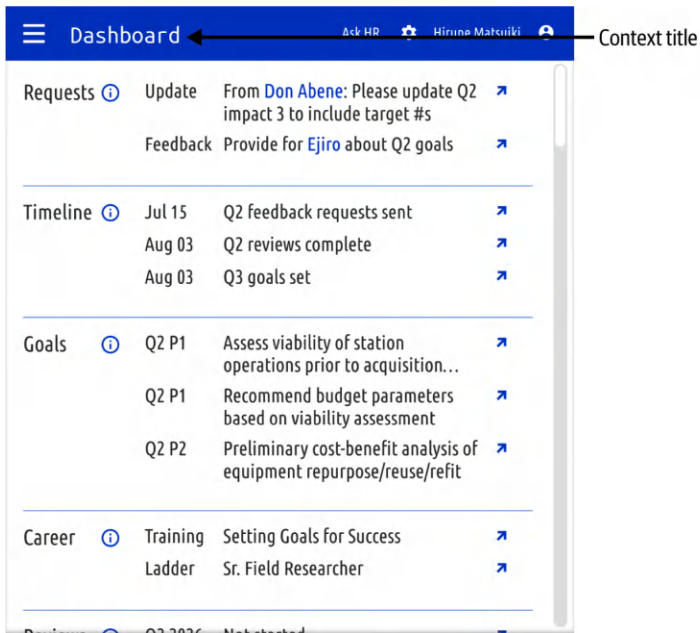


FIGURE 4-3

The context title of the LitMop home screen reinforces its role as a dashboard. People can use this dashboard to find tasks to take action on.

SINGLE-TASK TITLE

Single-task titles act as instructions for the person to take an action. Use an imperative verb phrase as a direction to reinforce the correct action.

For example, in Figure 4-4, the TAPP experience displays the code a person must scan when they get on the bus, to pay for their ride. The TAPP experience uses the single-task title “Pay Fare.” There is no button to use to take that action; the person must take the next action using the code-reader on the bus.



FIGURE 4-4

The Pay Fare screen in the TAPP experience has a single-task title in the app bar.

In Table 4-1, I summarize the guidance for titles. Throughout this chapter, I'll use curly brackets to show the text to be replaced. For example, {context name} is replaced in the examples with "The Sturgeon Club," "bluster," and "Dashboard." When a word is in straight brackets, it can be omitted. For instance, the guidance for single-action titles, the {verb} [the] {noun} pattern is shown in TAPP as "Pay Fare."

TABLE 4-1. Title pattern

Purpose	Provide high-profile clarity of context and the action to be taken.
Patterns	Context title: {Context name}; <i>for example, "The Sturgeon Club," "Bluster," and "Dashboard"</i> Single-task title: {Verb} [the] {noun}; <i>for example, "Pay Fare"</i>
Tips	Build confidence by using the words the person most expects to see. Keep titles brief so that keywords are noticed and recognized.

Titles, however they are used, are just the starting point. The real action takes place when people interact with the text by tapping, clicking, or otherwise selecting it.

As the single-task title demonstrates, the title's main purpose is frequently to introduce the action that the person can take. Most of the time, they will take that action by using a button.

Buttons and Menus

Buttons and menu items are the text or icons that a person interacts with by tapping, clicking, or speaking to get to their next step. Sometimes they are called *calls to action* or *commands*, but whether they are performing an action, taking the person to the next screen, or navigating elsewhere, we will consider them together in this pattern.

Buttons are some of the most important text in an experience. They are how the person makes their purpose known. Buttons (and to a limited extent, controls) are how people “speak” to the experience. The button must be used to enable the conversation between the person and the experience. Almost every other piece of text, including title, description, empty state, label, confirmation, error message, and more, is the experience speaking to the person.

Buttons and menu items work best when they are recognizable, specific, and only one or two words long. I have seen proprietary research that showed that buttons that are one or two words long were more frequently used than buttons longer than two words. Similarly, buttons that used a word that the person would actually say in a conversation outperformed generic buttons and buttons with words the person wouldn't have chosen.

For example, when someone in The Sturgeon Club wants to review and pay their club charges, they encounter the Folio screen (Figure 4-5). The button “Pay \$308.48” is the button most important to the club because this is its revenue stream. The text is a clear verb-first indicator of action. Putting the amount to pay on the same button makes it even more specific. Because the member already has a payment method registered in the experience, the payment can be completed with this single, seamless action.

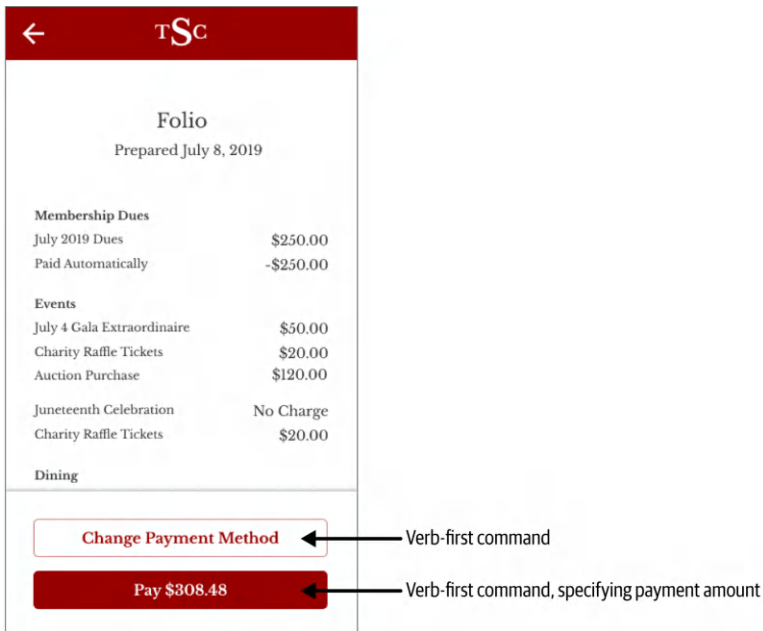


FIGURE 4-5

When reviewing and paying their charges in The Sturgeon Club, the person can use either the button Pay \$308.48 or the button Change Payment Method, or use the Back button to leave this screen.

There are two other options available to The Sturgeon Club member: the verb-first command of Change Payment Method and the back arrow. The order of the buttons is important: just like they would be in a conversation, the most common or primary action would be brought up first.

Sometimes, icons are used without words. Using an icon can also help reduce the number of words visible on screen. In the LitMop dashboard (back to Figure 4-3), arrows beside each item allow the person to navigate to take action. Screen readers used by people with low vision or blindness will mention the availability of these buttons, and they require text to be read for each of them.

This text that will be spoken by screen readers needs to be designed, even when it's not visible on the screen. There's an additional complication: sometimes the text is dynamic, and the designer won't know what content will be there in the future. In Figure 4-6, an excerpt from the LitMop dashboard designs shows the designer specifying the rules for

how the screen text would work, even though they can't know exactly what content will be displayed. Instead, the designer works with what they do know: for something to show up in those lists, the content will have to be specified and structured to have some predictability.

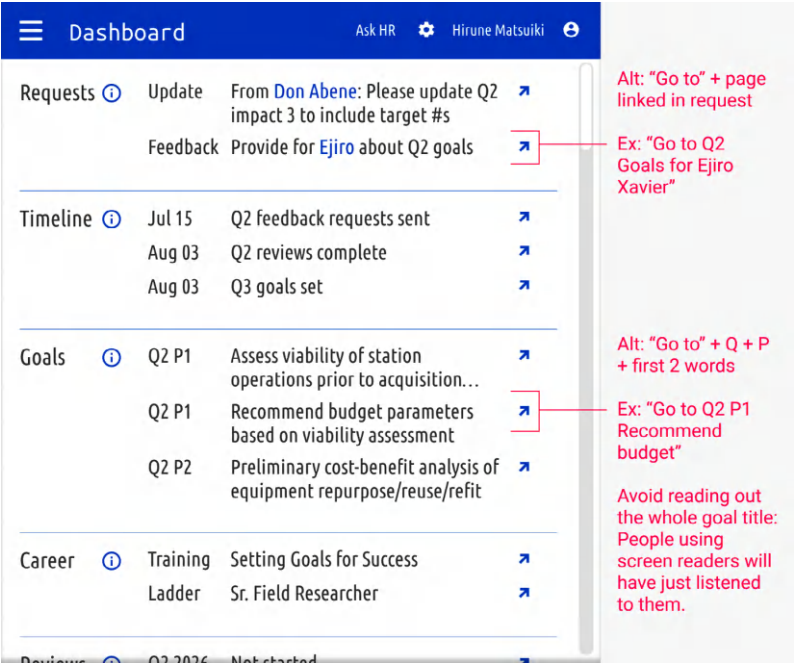


FIGURE 4-6

The LitMop design file includes annotation for the developers to use, specifying the content for the non-text screen elements that require text description, including these blue arrows on the dashboard.

In general, the same rules apply to these unseen words: the button will perform best if it maximizes clarity, is brief, and uses the same words the person would use in conversation.

Another type of button is an option in a menu or list. For these, nouns are sometimes more appropriate. For example, the 'appee menu gives the player access to the images they saved, their friends, the settings for the experience, and help (Figure 4-7).

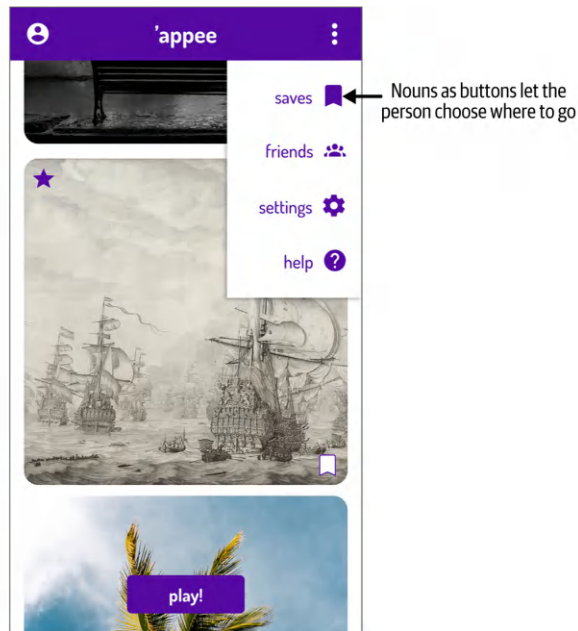


FIGURE 4-7

The 'appee menu is the way to access the player's saves, friends, settings, and help, which are examples of noun-based buttons in a menu or list.

Each menu item can be considered as its own button, but they all need to be designed in context, the way they will be used. This way, the words can be selected to differ widely from one another so that they are easy to disambiguate.

As 'appee shows, using only one or two words for each menu item enables the player to scan the options at a glance. In more complex experiences, it's even more important to select menu items as a set of options, together. An organization may need to consider entire taxonomies of menu items for shopping experiences or complex data storage, for example, so that the terms used are unambiguous throughout the experience. This design consideration will make it easier for people to choose the right option, the first time, every time.

When paired with single-action titles, buttons are most effective when they match the words in the title. For example, when a person needs to create an account in TAPP, they encounter the title "Create an Account" (Figure 4-8). The button that lets the person take that action, labeled "Create Account," matches the title. Because these two

phrases match so closely, there is no ambiguity: the person is committing the single action specified by the title. If the button said “Save” or “Submit,” it would be less clear to the person that they were taking the intended action.

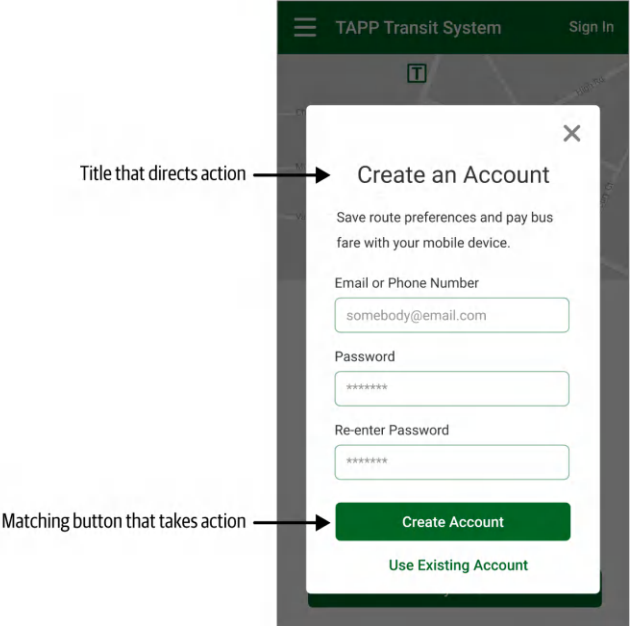


FIGURE 4-8 When someone needs to create a new account in TAPP, the title “Create an Account” intentionally matches the button “Create Account.” That symmetry reinforces the action to be taken.

Button text (whether in menus or not) is one of the most critical parts of the interaction design. These buttons determine how people think about the action they take; when we choose words they wouldn’t use or don’t understand, we create cognitive load that requires much more mental energy from them and increases their mistakes. When we get the text right (Table 4-2), we not only make it easier for the person using it, but we can also increase positive engagement and reduce costly support calls.

TABLE 4-2. Button and menu pattern

Purpose	Allow the person to advance toward or commit to action.
Patterns	Button: {Verb} [the {noun}] or {Verb} now, for example, “Play!” or “Create account” Menu: Group of buttons that works as a set, whether nouns or verbs
Tips	In English: 1-2 words, 3 words possible, for example, “Change payment method.” Use the words the person would say or write. For non-word icons or buttons, plan appropriate screen-reader text. When options will appear as sets (in menus, button groups), design them as sets.

In many cases, the button and title aren’t enough by themselves—the person might need additional information to be successful. They could require a reminder of the value of taking the action, and the space in the title and button is limited. To set expectations about how the experience will behave, or just to reinforce the brand, the experience can use descriptions.

Descriptions

We use descriptions to help people know what to expect, establish the brand, and reduce liability while they move forward using the experience.

A description is an informational chunk of text, sometimes called *body text*. Descriptions can appear as phrases, sentences, or paragraphs. A description can also be hidden text so that screen readers, such as those used by people with reduced vision, can explain a heavily visual design.

Most descriptive text can meet its purposes only if people read it. Descriptions are frequently ignored. Some team members despise them as “a wall of text.” *People don’t come to most experiences to read the UX text.*

Let's consider who does read the description: first, it's people who lack confidence in the experience or in their own ability to use it correctly. Perhaps it's a high-stakes experience, like applying for refugee status or setting up networked storage at a workplace. When it's technical or critical to get it right the first time, people read the descriptions.

The second group for whom the description is necessary are the people protecting the organization from liability. It may be important in the future, or when explaining the experience to regulators, to show how clearly important risks or other disclosures were made. This can be as ubiquitous as "All sales final" or "Your support session may be recorded for quality assurance purposes."

The third group is people who aren't fluent enough readers of that language to scan it briefly. These people will read every word, especially if the visual hierarchy isn't clear, because they can't tell what parts they can safely ignore.

The final (and smallest) group of people are those for whom it is a habit to fully and deliberately use any experience. This is less common among "digital natives" who have grown up using apps—sometimes even before they could read. Some folks will just read all the text that appears on the screen. Their experience would be incomplete without it.

As UX professionals, we need to advocate for all of these groups. The first group affects basic usage: people who could find it prohibitively difficult to use the experience without adequate description. The second group is represented by our legal advisors, product owners, and more. The text they draft may be discouraging for people in the first group, but we can help make it more understandable and palatable. For the third group, we can recommend localized experiences and the use of simple grammar and common vocabulary. For the last group, we need to remember to keep it all short!

HOW SHORT IS SHORT?

When description text is necessary, we should make it as easy to use as possible. In research shared with me inside a company, I learned that in English, people using an interface will rapidly scan lines up to about 50 characters wide, which is enough space for about three to six words. Similarly, people's eyes will linger on a few of the words

when a paragraph of text has three lines or fewer. Those few words are the description's opportunity to catch their attention long enough for understanding to develop.

When a single chunk of text becomes longer than these rough guidelines, people's eyes stop lingering on the individual words. They begin to feel more doubtful about it. Research participants and teammates will begin to remark about the "wall of text." Keeping the text brief and separating ideas into scannable chunks, no larger than the sizes previously described, makes people feel more confident about their own understanding and capability to use the experience successfully.

PLACE KEYWORDS STRATEGICALLY

To convey information to people who will barely glance at the text, it's essential to place meaningful words where they'll be noticed. First, this means understanding which words matter to them. These can be the words that they'd be anxious about seeing, like "theft" inside a financial app or "toxic" in a medical app. It can also be words that reassure them that they're in the right place, to create a sense of confidence.

While we might include these words in the title, we usually don't have enough room there. When we're using them in the description, it can be most effective to put them at the end of the phrase. For people reading English, left to right, they tend to scan in a left-to-right, top-to-bottom pattern. This "F-shape" scanning behavior means that words at the ragged end of sentences are more likely to be noticed. The second most effective place to put keywords is at the beginning of sentences, but because English text is frequently left-aligned, that's also the most crowded place for the eye to be taking in information.

USUALLY AVOID ASTERISKS (*)

Trust is essential. We reduce trust when we give people beautiful promises and easy paths forward, but take those promises away with asterisks and fine print. When we make clarifications and disclosures difficult to read, we communicate that we're willing to work to hide things from them. Using asterisks can indicate that the main text isn't fully honest and can't be trusted.

If there are complex ideas that must be included while somebody is using the experience, include those ideas in descriptions. Use plain language, and include how it benefits the organization and the person

using the experience, as necessary. This usually requires collaborating closely with product owners, attorneys, privacy professionals, and business owners.

For example, The Sturgeon Club includes a messaging system within its experience, as shown in Figure 4-9. The Sturgeon Club uses a description at the end of the message list that sets expectations for members that the message system is secure, that it works only among club members, and that messages are deleted after 30 days. Although it can be useful information to the member, and certainly useful for The Sturgeon Club to have told them, the member doesn't need to read the text to successfully use the messaging system.

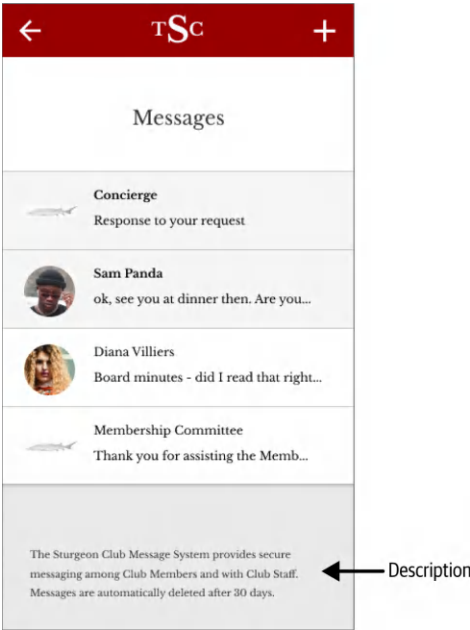


FIGURE 4-9

This Sturgeon Club Messages screen has two read and two unread message threads. The description at the bottom of the screen sets member expectations about how the messaging system works and indicates that messages will automatically be deleted after 30 days.

When a player wants to play, 'appee uses descriptive text on its “basic rules” page (Figure 4-10). The person doesn't need to read the description to be successful, unless they were planning on breaking the rules. The purpose of this descriptive text is to remind the player of the rules

and to force acknowledgment of those rules before the person can play. This is a step that 'appee takes to reduce its liability for inappropriate images or claims of unfair votes or decisions.

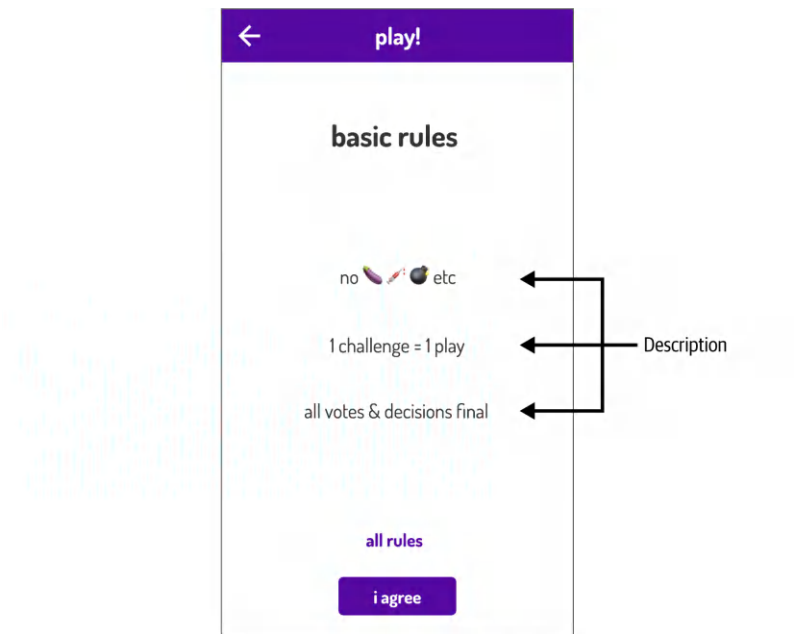


FIGURE 4-10

The basic rules screen in 'appee mostly contains descriptive text to set expectations that inappropriate sexual references, drug references, and violent topics are not allowed; that the person can play only once per challenge; and that all votes and decisions are final. This screen mostly exists to protect 'appee from claims that their rules are unavailable or unremembered.

To satisfy people who want all of the rules, who will be reassured by that wall of text, or who need specific information to continue, 'appee includes an “all rules” button. This is one way to include additional information as an option instead of as a requirement to continue.

We can use word order in our sentences to emphasize different things. For example, we could emphasize that the person should take an action with *{Verbing} [the {noun}] helps you do X*. If we needed to explain something that already happened, we might choose *If X happens, you'll see Y*. In contrast, if we want to emphasize that some desired outcome is available to them, we might choose *To do X, {verb} [the {noun}]*, as we see in this example from TAPP (Figure 4-11).

The first screen in TAPP’s “Buy Fare” process includes various passes that can be purchased by the person currently signed in, but if the person is eligible for a reduced fare, the description provides a different path forward. The rider can continue and purchase it at the regular price, but if they read the description, and they are eligible, they can take the steps described to get the reduced fare.

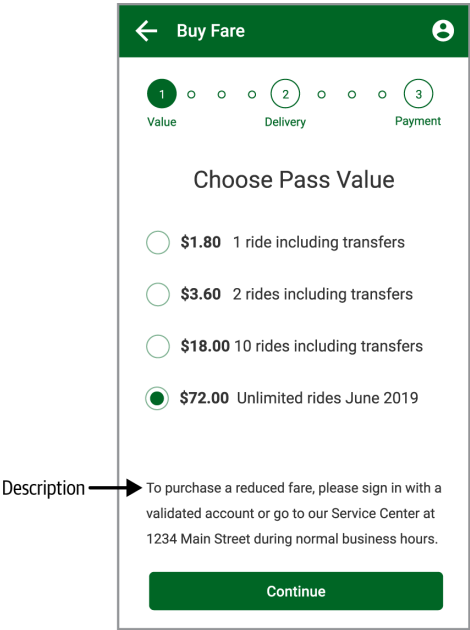


FIGURE 4-11

The first screen to buy bus fare in the TAPP experience. The description tells people that to purchase a reduced fare, they can sign in with a “validated account” or go to the physical location of the TAPP Service Center.

Note, this description might feel incomplete. For example, there is no information about what a reduced fare is, or what a validated account is, or how one might know whether they qualify. To keep the text brief, scannable, and on point, that information must be optionally available. The job of this screen is to allow the rider to purchase a fare, not to become informed about the fare. The TAPP experience must provide enough information elsewhere in the experience for people to make that choice.

LitMop attempts to provide contextual information using tool tips. This common interaction model allows additional information to be available if people want it: they can interact with an icon or other affordance, like dashed-underlined words, to reveal more information. Getting these to work correctly with screen readers can be challenging, but it's possible.

In the Dashboard, each section has an information symbol (i). When a person selects it, a tool tip is displayed that can have a title, text, and linked text. In the one displayed in Figure 4-12, the person is shown what kind of tasks they can expect to see in this section of their dashboard.

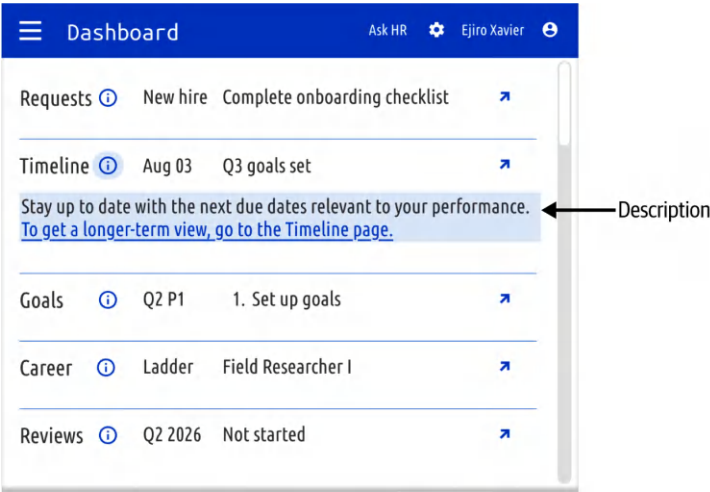


FIGURE 4-12

The tool tip for the Career section is open on the dashboard for a new employee, which contains a title and description.

In the Timeline tool tip, the description also contains a link to a different place within the experience. This isn't always a great idea, because you may be distracting the person from the task at hand. But in this dashboard context, the whole point is to get people to the pages where they can complete their tasks.

I put this example in because I want to point out this best practice: when providing links (instead of buttons), link the entire sentence or meaningful phrase, not just “timeline page” or “page.” A common way to use a screen reader is to move between links so that a person doesn’t have to listen to a lot of exposition and repetition before getting to the thing they want to use. For the same reason, we avoid using “click here” as the place to anchor a link.

The description pattern (Table 4-3) is one of the most flexible. Remember, though, I’m not prescribing any one of these patterns. Instead, I’m trying to provide starting places for your own iteration, to help you start with something usable even before you tune and tweak it to best work in your experience.

TABLE 4-3. Description pattern

Purpose	Set expectations, establish brand, reduce liability.
Patterns	To do X, {verb} the {noun}, for example, “To purchase a reduced fare, sign in with a validated account...” {Verbing} the {noun} helps you do X. If X happens, you’ll see Y.
Tips	In English, use 3 lines max, 40–50 characters per line. Usually avoid asterisks. Use keywords at end of lines. When adding hyperlinks, link the entire phrase.

One of the most elegant places for titles, descriptions, and buttons to work together is in an “empty state” in the experience. When the expected action or content isn’t available to the person, the experience can feel empty. Let’s examine this special case of titles, descriptions, and buttons next.

Empty States

We use empty state text to set expectations and build excitement while indicating that the empty space is intentional. Examples of empty states include shopping carts with nothing in them, inboxes with no emails or messages in them, and an area for saved games before your first play.

In my experience, product teams tend to start designing any feature with the optimal case: a person is fully engaged in the experience, using it to its best potential. Experiences that are designed to highlight the things the person has done or acquired can feel pretty empty when the person first opens it and nothing is there. Content designers can use empty state text to indicate that the emptiness is not a mistake.

Empty state text can be as simple as a single line of text, or as complex as a title, description, and button.

For example, when a member isn't signed in to The Sturgeon Club, no other action can be taken (Figure 4-13). The only way forward is to sign in. Even the menu is emptied to prevent people from taking any action. The empty state text in the menu, therefore, helps members move forward: “To access your membership, [sign in](#).” The “sign in” portion is interactive text that will begin the sign-in experience.

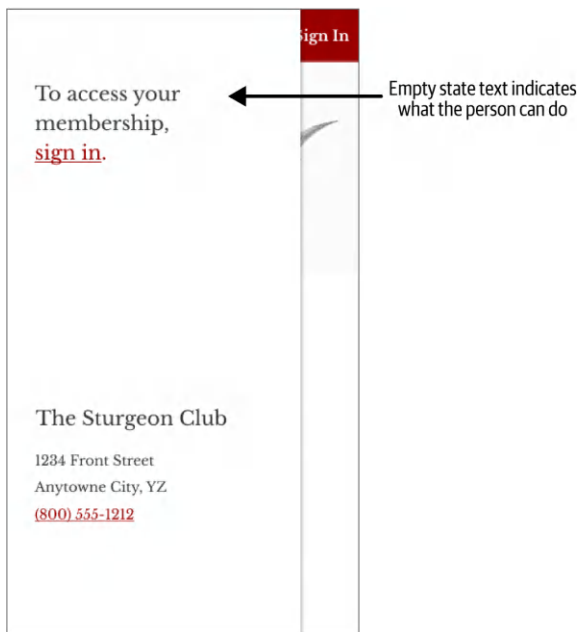


FIGURE 4-13

When nobody is signed in to The Sturgeon Club, empty state text directs people to sign in.

In some empty state conditions, there is nothing the person can do to fill it. In 'appee, if you open another person's profile and that other person has never submitted an image, there are no images to display. For example, the profile page for goldilox has no images to display because they've never submitted an image (Figure 4-14). The text tells the viewer "when goldilox plays, see their entries here."

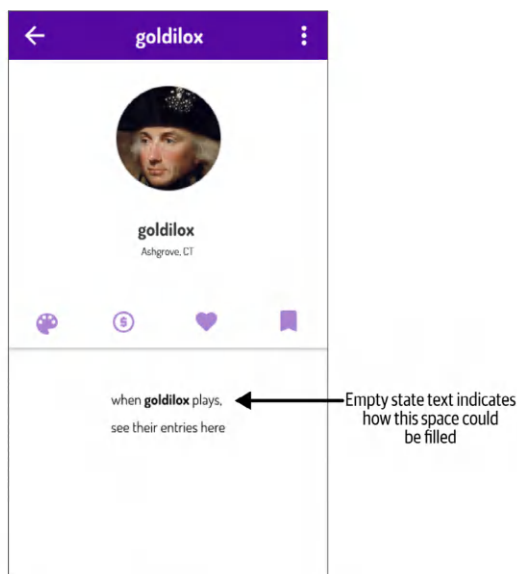


FIGURE 4-14

The view of a person's profile in 'appee usually displays the images they have entered for different challenges. When that person hasn't entered any images, other people viewing their page will understand how that space would be filled.

Empty states can get more complicated. For example, it might be impossible to give the person a one-step action to take, but they might be very interested in filling it. In TAPP, it's very handy to save a commonly used bus route. But before any bus routes are saved, the screen has no information to show (Figure 4-15).

Instead of displaying simply "No routes saved," TAPP uses the opportunity to educate the person about how to save those routes. The empty state provides the instruction "To save a route, tap Save when you find

the route you want.” Then, it provides a button to find a route. Together, the instructions and the action move the person toward successfully saving a route.

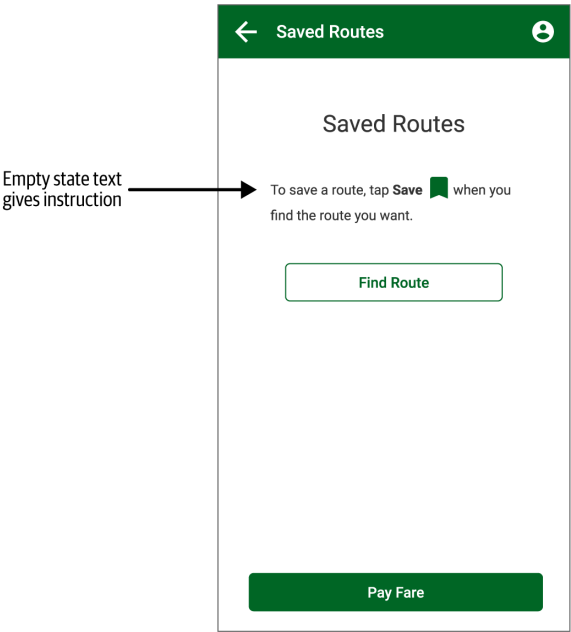


FIGURE 4-15

When no routes are saved, TAPP instructs the person on how to save a route. Then, it provides them a way to take the first step in following those instructions: finding a route.

Empty states can be found within forms, too. In the Add Goal example from LitMop, the LLM doesn’t yet have enough information to draft the success report for that goal. To keep the experience consistent, the person is shown a disabled state for the success report editing field (Figure 4-16), with text explaining, “When a draft report is generated, you’ll be able to edit it here.”

While I provide patterns as a starting place (Table 4-4), these otherwise empty moments can be used to showcase your brand. It’s a great place to exercise your brand voice, iterating with concepts and vocabulary to highlight the uniqueness of your product, company, or feature.

Goals Ask HR Hirune Matsuiki

Measured by ⓘ Target Unit

Select a metric Enter # Select unit

Scope ⓘ

Example: During onboarding in the Acme Product, among users within North America

Success report ⓘ

When a draft report is generated, you'll be able to edit it here

Edit for accuracy. The best reports tell how success was accomplished and include relevant project names and numbers.

Empty state text sets expectations for what the app will do

FIGURE 4-16

In the LitMop Add Goal experience, the success report field is present but disabled, but it includes empty state text that helps the person understand how it will get populated and how they'll be able to interact with it in the future.

TABLE 4-4. Empty state pattern

Purpose	Set expectation, build excitement, indicate intentionality.
Patterns	<p>Title: {Verb} the {noun}</p> <p>Description: To do X, do Y</p> <p>Button: {Verb}</p> <p>Or if no action is available:</p> <p>Title: {Area name} or omit</p> <p>Description: "If/when X happens, you'll see Y," for example, "When a draft report is generated, you'll see it here."</p>
Tips	Great place to showcase brand through information and attitude, using voice.

Descriptions, buttons, and titles are great tools for empty states. But when the experience is the opposite of empty—when it’s full of rich detail—we need a specialized form of description: labels.

Labels

Labels minimize the effort required to understand the experience with noun phrases and adjectives that name or describe things. They are used to indicate sections, categories, status, progress, quantity, or unit. Labels are ubiquitous where there is a lot of detail to convey, because they communicate a lot of information in a compact, no-nonsense format. But even labels deserve attention: they still need to be chosen, align to voice, and be translated and/or internationalized.

The difference between a description and a label, in these text patterns, is a matter of length and purpose. Descriptions are typically full sentences, regardless of punctuation. Labels are usually single nouns or noun pairs. Descriptions are frequently used with titles, buttons, or whole experiences. Labels are generally related to passive screen elements like icons or sections and limited to that local context.

To meet their purpose, labels must use specific terms and avoid unfamiliar jargon. If vague terms are used, or jargon that the person doesn’t understand, the labels can increase the effort to understand the experience. Labels are an important place to focus usability testing and other user research to uncover the words a person would naturally use for these labels. The words that are already in their brain will be the easiest ones for people to read and understand.

Labels are frequently complicated by including dynamic elements. That is, the designer won’t know ahead of time what the cost of an item might be, or the date, or the number of “likes” a social media post might garner. To be successful, they’ll need to know the variables they’re working with and choose words that will work with all possible values of those variables.

For example, on the Folio screen of The Sturgeon Club, the labels include the date label, section labels, the monetary labels that indicate the cost of each item, and the text label “No Charge” (Figure 4-17). When this Folio experience is created, the date label could be written as “Prepared {date},” where {date} represents the date that the folio was prepared for the club member. The format of the date needs

to be further specified as the name of the month, then the number of the day, then a comma, and then the year. It's often up to the content designer to specify that date format, in partnership with design and engineering.

Engineers should use existing code libraries, when available, for number formats like dates and money, but those labels should be checked. For The Sturgeon Club Folio, the decision had to be made to use a minus sign, instead of putting the dollar amount in parentheses. Perhaps members were asked, in a research scenario, how they would expect adjustments or payments to be indicated. The Sturgeon Club charges in US dollars exclusively, so it makes sense to use “\$” and “.” to separate dollars and cents, but an experience also available in Europe, for example, would need to consider how it would use “€” and “,” to correctly represent the numbers.

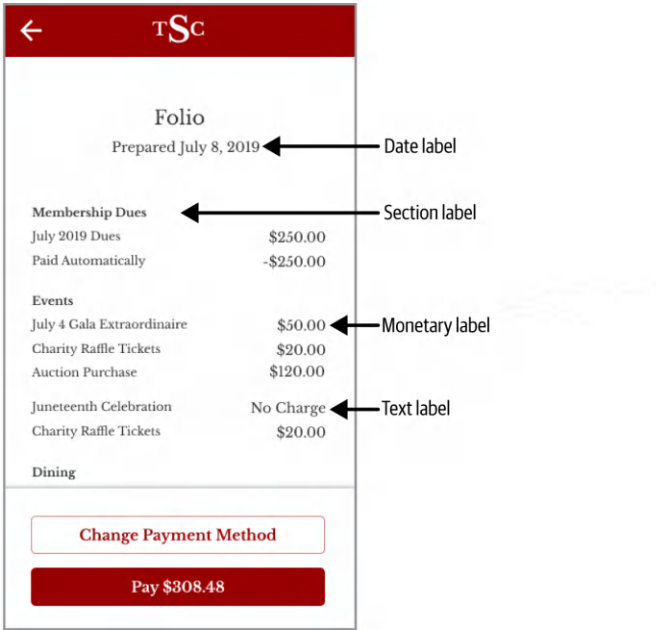


FIGURE 4-17

The Sturgeon Club Folio screen uses labels to indicate date, sections, currency and quantity of money, and when there is no charge.

In the 'appee experience, when a person views an image, they can also see several labels that show statistics about that image (Figure 4-18). A row of icons act as buttons, even though no visible button text is provided. This is usable only if the icons are instantly recognizable, but it is aligned with 'appee's voice of using fewer words than strictly necessary. Instead, the labels provide context: the person can see (or hear, using a screen reader) that more than 12,000 other people have left comments, 238 people have made purchases related to this image, more than 102,000 people have "liked" the image, and more than 382,000 people have saved it. The 'appee decision to use a lowercase "k+" to indicate "more than...thousand" is another reflection of its voice; it uses sentence case when providing insights.

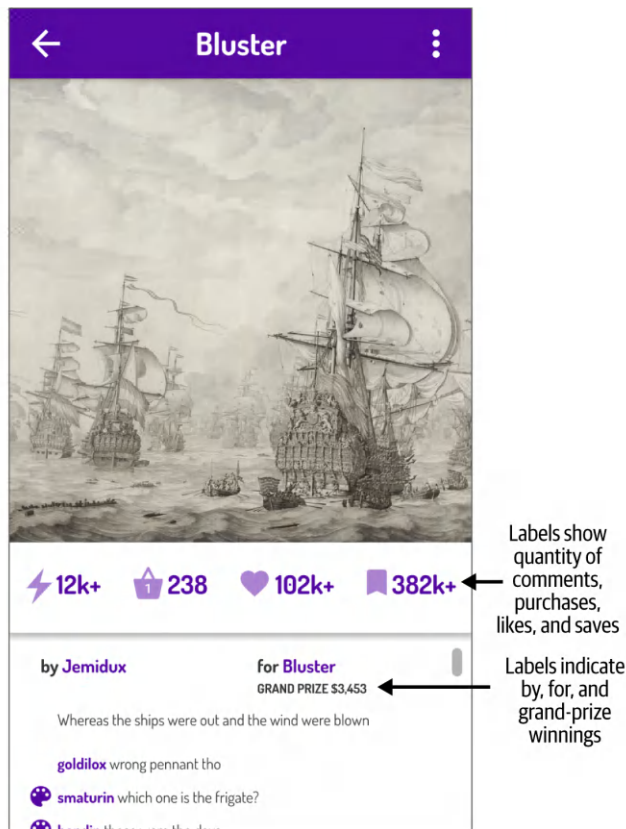


FIGURE 4-18

Labels on the 'appee "view image" screen indicate the quantity of comments, purchases, likes, saves, and winnings of the image and differentiate between the player and challenge.

The 'appee labels “by” and “for,” which indicate the artist and challenge, respectively, are in danger of not localizing well, because not every language uses prepositions the way English does. The content creator should work with their internationalization expert and designer to create alternate layouts, in case longer words are needed to convey the same meaning in other languages. For example, the alternate label “artist” or “player” could be used instead of “by,” if it could be placed vertically above the player’s name.

We also see labels in LitMop. When a people manager is using LitMop to check on and report their teams’ goals, they can review all of the inputs into the reports in a streamlined table. Each column of the table is labeled: the due date, the metric being measured, the target and units for that metric, and the project and region related to that goal (Figure 4-19). As a whole, the labels help the manager make sense of the information provided.

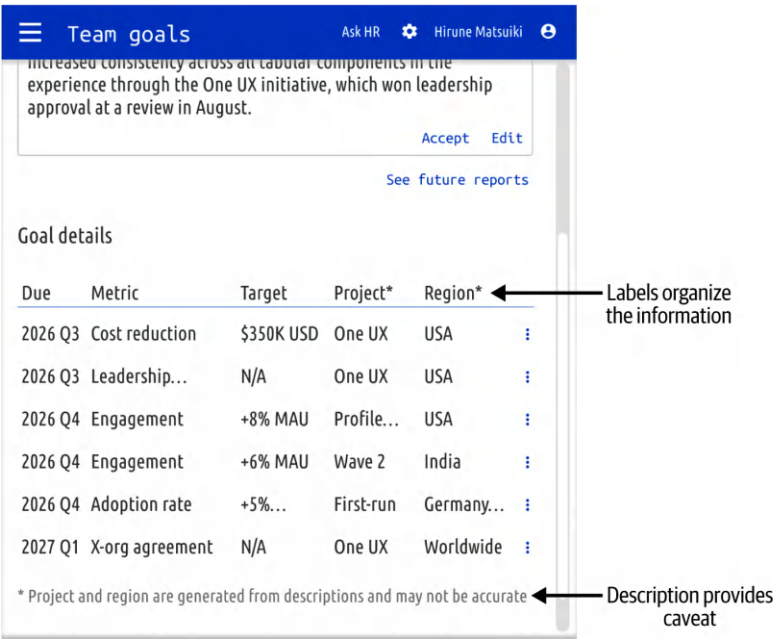


FIGURE 4-19

The “Team goals” page in LitMop shows a table of all goal details, with labels for table headers. Asterisks on two labels call out the potential inaccuracy of the information in those columns, as described in a footnote.

In this screenshot, also note a special case for asterisks! LitMop needs to convey that “Project” and “Region” are identified by the embedded LLM and may be inaccurate. The column format doesn’t allow space for that information beside the label, so the asterisk is appropriate. While they could have hidden the information in a tool tip, LitMop reduces its liability for inaccurate information by making that caveat visible with no interaction necessary.

It’s difficult to call out a pattern for labels (Table 4-5) other than by using single words, units, and numbers as appropriate. There are so many possible words, or sets of words, that can be appropriate labels, in the diverse experiences you create.

TABLE 4-5. Label pattern

Purpose	Convey accurate information while minimizing the effort required to understand the experience.
Patterns	{noun} {unit} {number}
Tips	Be as specific as possible. Ensure that labels are easily differentiated from one another within a set.

Labels are a specialized form of description, in that they tend to be briefer and more technical, but they are still distinct from buttons because people merely read them. Now let’s go one level deeper into specialization: the unique names and statuses for controls.

Controls

We use controls to allow people to make customizations and inform people of the extent and state of those settings.

To write for controls, I find it helpful to acknowledge that many of our controls use as their core metaphor the analog dials, switches, sliders, and indicator buttons of early electronics and mechanical devices. In most meaningful ways, the use of the physical and software switches is the same. At its best, the categories and labels make the possible range of customization clear.

There are usually at least two pieces of text to consider for any particular control: name and state. The name should be the noun or verb phrase that names or describes the control in a way the person using the experience will recognize. The state of the control is, for example, whether a checkbox is checked, the position of a slider, or whether a toggle is flipped left, right, up, or down.

The UX text that is paired with the control needs to match the possible states of the control. For example, a checkbox indicates the affirmative when checked, and the negative when not checked. If we don't choose a name that has a clear meaning in the affirmative and in the negative, it won't work with its checkbox.

The setting state text can be visible or invisible, but a screen reader will still read it out loud. Checkboxes will be read as "checked" or "unchecked." Toggle switches have implied states of on or off, but might be labeled with a similarly opposite pair of states: red/green (with differentiation for people who are red/green colorblind), enable/disable, or other. Sliders and dials can use state text to establish the endpoints of the range or use the implied text of maximum and minimum.

We also need to consider grouping controls together in a list. The Settings page from The Sturgeon Club demonstrates how much work a group name does to set context for the controls (Figure 4-20). Each control could be listed separately: "Show Today's Events on home screen," "Show New Messages on home screen," and so on. But grouping them together makes the whole list easier to understand, with less reading. This parallel construction makes it easier for the reader to understand not only each list item, but also how the whole set of items works together.

Another consideration for the name of a control is how the person will be directed to it when they need help. For this reason, having a unique name for each control can be important, even if those controls appear in different sections of the same page. Category names can be verb phrases, like "Show on home screen" in The Sturgeon Club, and they can be noun phrases, like "Account" and "Notifications" on the TAPP settings page (Figure 4-21).

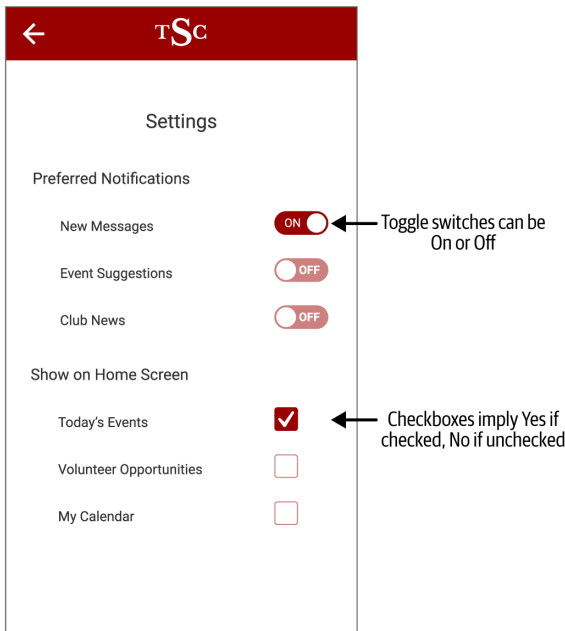


FIGURE 4-20

The Sturgeon Club settings allow members to switch notifications on or off by type. They can also choose which kind of content to show on their view of the home screen.

Complex design systems may have many different components for the controls available in a particular product or product suite. Embedding the control pattern (Table 4-6), or your own pattern, customized to that product, inside those components can guide subsequent designs to create more usable content from the beginning of the design process.

TABLE 4-6. Control pattern

Purpose	Allow people to make customizations and inform people of the extent and state of those settings.
Patterns	{Verb} [the] {noun}, for example, “Save payment info” {Common setting name}, for example, “Notifications”
Tips	Group similar controls under the same heading to reduce text. Use parallel construction to reduce confusion.

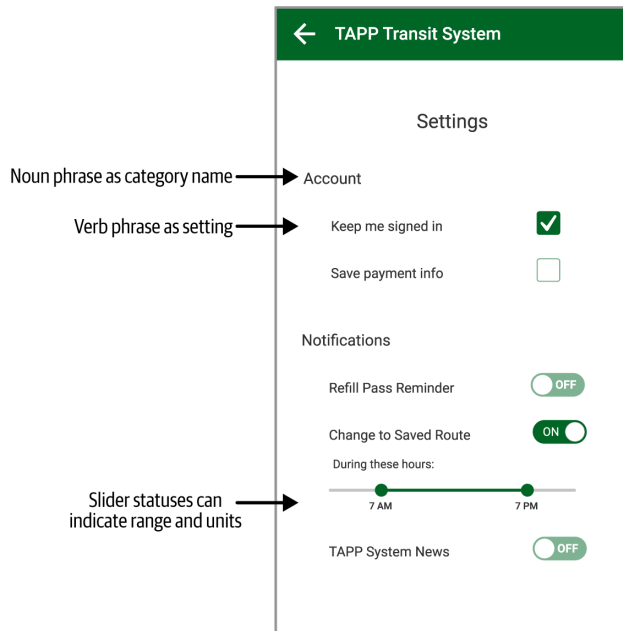


FIGURE 4-21

TAPP uses parallel construction at the category level, with nouns “Account” and “Notifications” setting context for those sections. The control names in the Account category are verb phrases that are related to the person’s account. In Notifications, each control is named with a descriptive phrase.

The usability of the control depends on the person recognizing the purpose of the control and understanding how that matches to their own needs. Similarly, people need to recognize what kind of text to enter in a text input field, whether it’s there for a message, number, password, or other kind of text. Next, we examine text field labels, hint text, and default text.

Text Input Fields

To help people enter accurate information in input fields, we use a variety of text surrounding those fields. Form fields use UX text as labels, hints, and prefilled text for entering text, email addresses, numbers, dates, and other information.

Using the person's information or LLM-generated text to prefill the text field can save the person time and give them the opportunity to correct it. But this works only if the experience already knows that information and knows that it is very likely to be correct.

When we can't prefill the text field, we use labels outside of the text field as well as hints within the text field to indicate what content the person should enter.

With hint text, we need to be careful. Some research indicates that people can interpret hints as prefilled text. When we do decide to offer hint text, the label and hint can work together to provide more guidance than either one could provide separately.

For either the label or hint text, there are four good options for the text:

- Name of the information to be entered
- Example of the information to be entered
- Verb-first instructions about entering information
- Guidance for how the person can be successful

Using these options consistently can help build a person's confidence that they are entering text correctly. *But even more important than consistency is clarity.* If you can make the path forward clearer and make people more successful, it's better to use inconsistent options for the UX text pattern on the same screen.

For example, the Change Password overlay in The Sturgeon Club experience contains a New Password entry field (Figure 4-22). The hint provides guidance of what will work in that field: "At least 8 numbers or letters." In the same screen, the other two text fields have their own pair of label and hint designed to support the success of the member.

← TSC ⚙

✕

Change Password

Current Password Names the info to be entered

***** Masking asterisks hide secret info

New Password Names the info to be entered

At least 8 numbers or letters Provides guidance

Confirm New Password Instructs what to enter

Same as new password Provides guidance

Save Change

Cancel

FIGURE 4-22

For The Sturgeon Club, when a person wants to change their password, they need to enter their current password, then their new password, and then their new password again to confirm it. The design uses pairs of labels and hints to help the member change their password successfully.

In some cases, a team might choose to use only hint text and avoid using labels. In many cases, this can help the design look clean and minimal. In those cases, the hint text must do all of the work of informing the person what text they should enter. When the person begins to enter information, there will be no label for the text box, so it's not a design pattern that maximizes usability.

For example, signing in to 'appee requires the person to use their phone number or email address and their password (Figure 4-23). The hints are the names of the information to be entered: "email or phone" and "password." 'appee relies on the player's recognition of this common pattern to successfully sign in. Note that 'appee is making a decision that is counter to best practices for usability: when the person begins entering information, there is no label to indicate what should go in either input field.

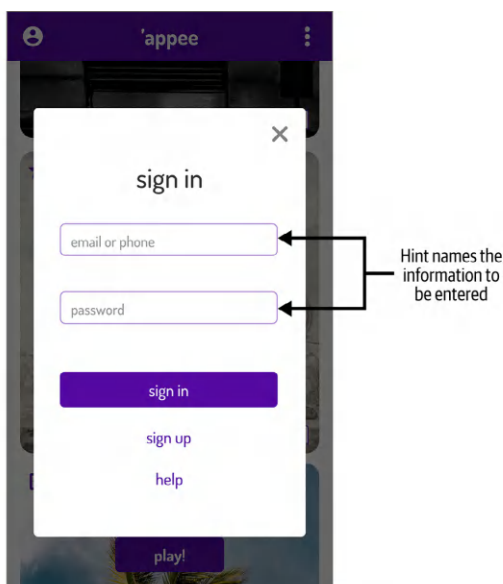


FIGURE 4-23

The 'appee sign-in screen uses the names of the information to be entered (email or phone along with password) in its two text fields.

The Request Help screen in the TAPP experience provides an example for default text field entries (Figure 4-24). TAPP has an email address for this signed-in TAPP rider, so their email address is entered for them by default when they come to the page. The person might enter a different email address, but they don't need to.

In LitMop, input fields in the Add Goal screen use a variety of hints (Figure 4-25). For Measured by, Target, and Unit, the person is directed to fill the field appropriately. In contrast, the Scope field gives an explicit example.

In the "Success report," the hint text reminds people about the criteria for the best possible reports. In this case, since this field will be filled by LitMop using an LLM, it's appropriate to remind the person that they still have an action to take: editing for accuracy.

As we see in Table 4-7, there are multiple general patterns for text input fields. These should be chosen primarily for clarity, and then iterated upon to maximize usability and brand appropriateness.

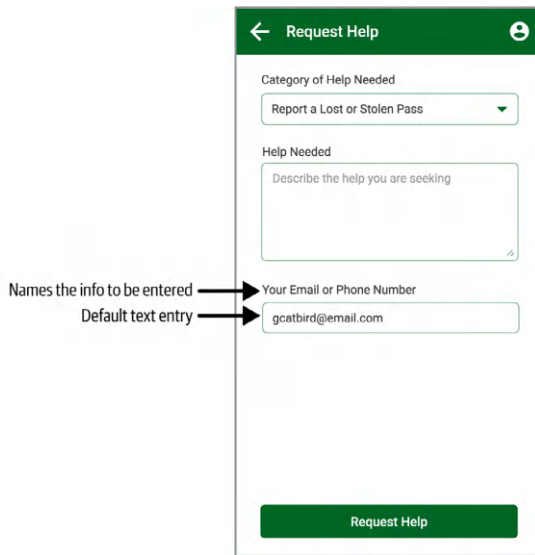


FIGURE 4-24

On the Request Help screen in TAPP, the email or phone number associated with a TAPP account appears in the text field by default.

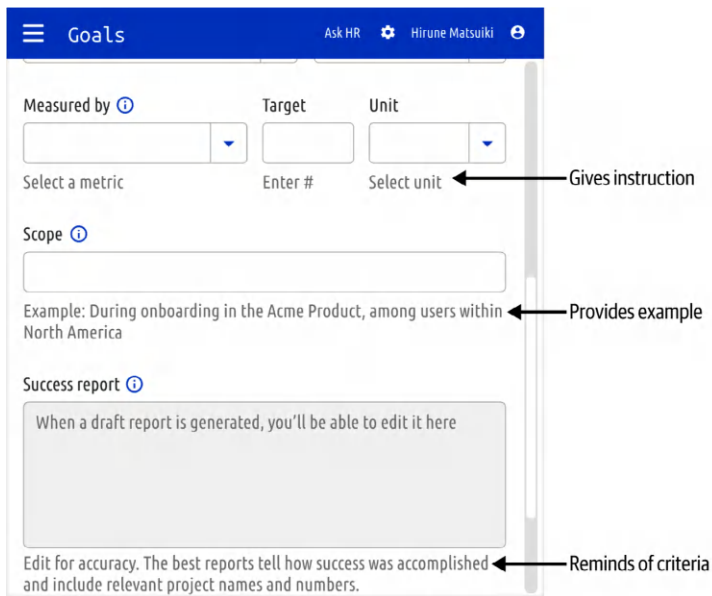


FIGURE 4-25

On the Add Goal screen, LitMop uses a variety of methods in the UX text to get the person to use the fields accurately, including prefilling the fields.

TABLE 4-7. Text input field pattern

Purpose	Help people enter accurate information.
Patterns	{Common information name}, for example, “Name” {Verb} [the] {noun}, for example, “Confirm new password” Example information, for example, “yourname@email.com” Guidance for how they can succeed, for example, “Describe the help you are seeking” (Prefill the appropriate text for them)
Tips	If hint text is used as labels, make sure sufficient context exists when fields are filled.

After a person has successfully used a set of text fields, there’s often a pause while the system transmits and/or validates those inputs. That pause can be stressful for the person, especially if the information they entered is sensitive, like an account number, or complex, like an online job application, or emotionally fraught, like a message sent to a romantic interest. A kind, easy step for the experience to take is to make that pause visible. Although a spinner or other animation is usually enough for a person with full vision, we can provide text on the screen or for the screen reader to help with that transition.

Transitional Text

We use transitional text to confirm that an action is happening, when there are few or no other indicators.

When an experience “hangs,” or is delayed while an action is processed, it’s courteous to inform the person that their waiting is not in vain. Just as a person at a help counter will say, “Just a moment while I get that for you,” a digital experience can use transitional text to indicate that it has received a request and that the person will need to wait for a moment.

In general, transitional text shouldn’t require any additional action or taps from the person. If the action is in progress, use the present continuous tense of the verb, like “is uploading” or “are sending.” Ellipses can help to indicate that the delay will be brief, especially if an animation isn’t present.

For example, after a person updates their payment method in The Sturgeon Club, they get an overlay message that the updating continues (Figure 4-26). They can feel confident that they have done what they

intended to do and that if they wait, the process will complete. Showing this message also prevents the club member from using the old payment method accidentally before the process is complete.

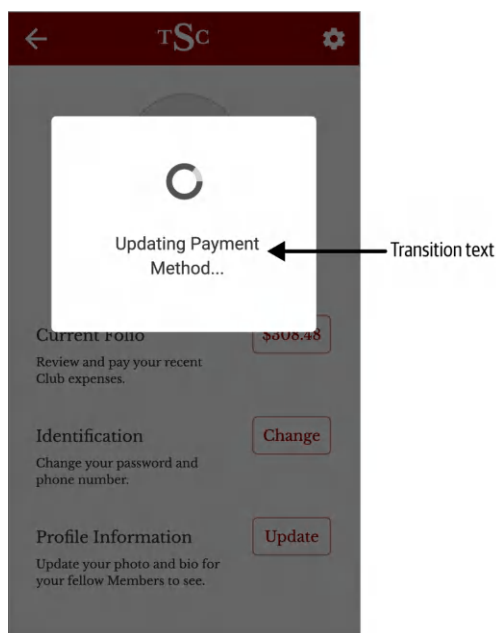


FIGURE 4-26

Transitional text in The Sturgeon Club helps build the club member's confidence that they have successfully done their part to update their payment information, while setting the expectation that the process of updating will be complete in a few moments.

If it's appropriate in the experience, we can heighten a sense of excitement by introducing a delay. In 'appee, it can take a few moments between when a person agrees to the rules and when the current challenge is retrieved from the service. That delay can even be extended to build anticipation further to ensure that even if the database responds quickly, the player has a few moments for anticipation to build. The transitional text in Figure 4-27 emphasizes that the challenge is being prepared and that there's "no turning back now."



FIGURE 4-27

When a player agrees to the rules, there is a delay while 'appee retrieves the current challenge and registers this person for this challenge. The transitional text helps to build excitement with the phrase “no turning back now ~,” even as it confirms that progress is happening with the more common phrase “getting ready ~.”

Generic verbs like “getting” work well for many ambiguous circumstances, but in most cases, specific words work better. When a rider taps a Map button for a particular route, they expect a map of the route and possibly details about its closest stop. When TAPP experiences a delay in retrieving those details, it still demonstrates it received the Map button press by displaying the transition text (Figure 4-28).

When a person adds a goal to LitMop, there may be a delay as the success report is generated (Figure 4-29). That information is conveyed in a state change in the “Success report” section, including a change from the placeholder text (previously shown in Figure 4-16).

After the transitional text, the changing of the experience can be its own confirmation, such as when LitMop’s LLM-generated text appears (Figure 4-30).

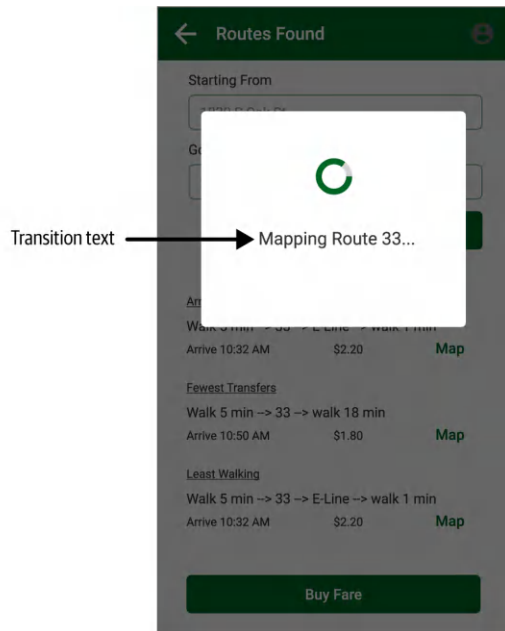


FIGURE 4-28

When TAPP has a delay retrieving a route map and data, it displays the transitional text “Mapping Route {number}...” to build confidence that it is working on performing the action that the person requested.

Most examples of the transition pattern rely on the verb tense that means “this action is still happening.” For circumstances where the verb tense, by itself, doesn’t work, it can help to say that {the name of the process} continues (Table 4-8).

TABLE 4-8. Transition pattern

Purpose	Confirm that an action is happening.
Patterns	In English (and other languages with verb tense): {Verb}ing [the {noun}][...] {Process name} continues[...], for example, “Updating payment method...”
Tips	Use ellipses (...) to emphasize that the transition will be brief and that the person should wait.

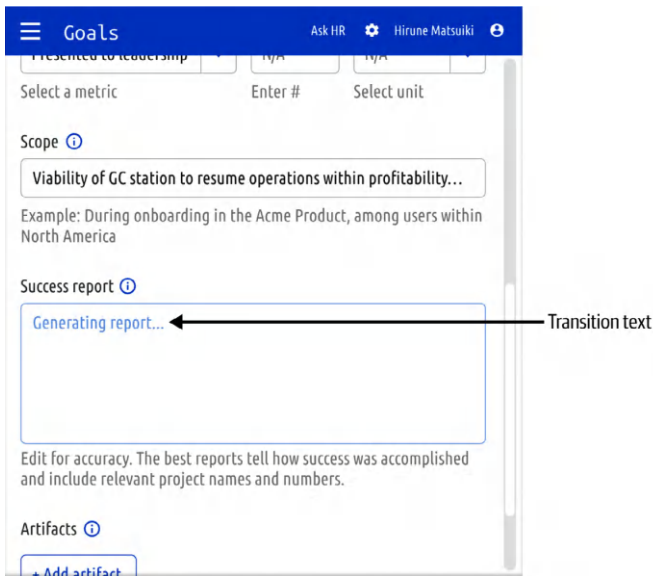


FIGURE 4-29

When the LitMop LLM is working, the state of the input field changes and the text changes to “Generating report...”

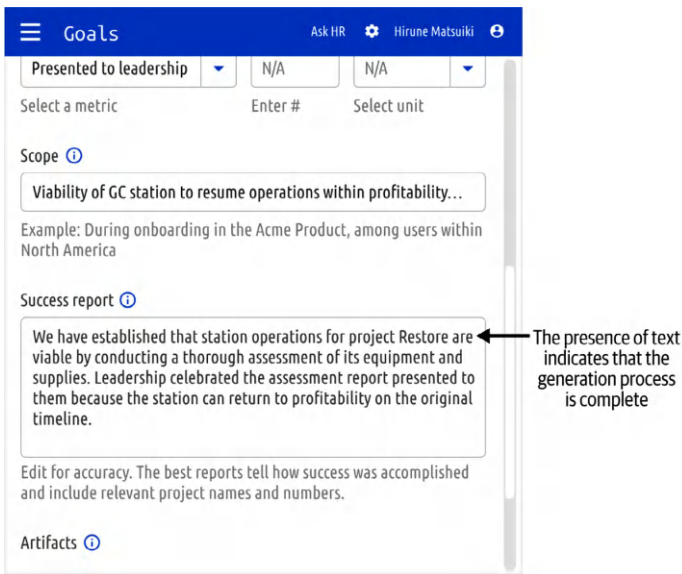


FIGURE 4-30

When LitMop’s generation is complete, the appearance of the generated text is all that is needed, without requiring a confirmation message.

When the transition creates a clear change in the experience, we sometimes don't need additional text to emphasize the point. But when the effect is more subtle, it's a good idea to provide a confirmation message.

Confirmation Messages

We use confirmation messages to inform or reassure the person that the progress or results they expect are complete.

These are especially useful when the result of an action is delayed. These confirmations can appear passively in the experience while the person continues forward, or they can appear in sequence in the experience as a momentary pause or step.

The basic pattern for confirmation messages is to use the past tense of the verb or verb phrase that best describes the action. In English, using one verb in the present continuous tense for transition (submitting) and the same verb with the past tense for the confirmation (submitted) can provide a sense of completion. Similar verb pairs include sending/sent, removing/removed, deleting/deleted, and posting/posted.

Confirmation messages allow the experience to continue while other systems work. For example, The Sturgeon Club provides the transitional text "Saving..." while a person is entering text and the confirmation "Draft Saved" when a person pauses as they enter a message in the secure messaging system (Figure 4-31). Similar messages can be seen in Google Docs and Microsoft Word Online when the document is being saved online in real time.

The confirmation text can be a single word when the context and action to be completed is the person's sole focus. When a person submits an image to 'appee, the transition text "submitting your entry ~" is replaced by the single word "submitted!" (Figure 4-32).

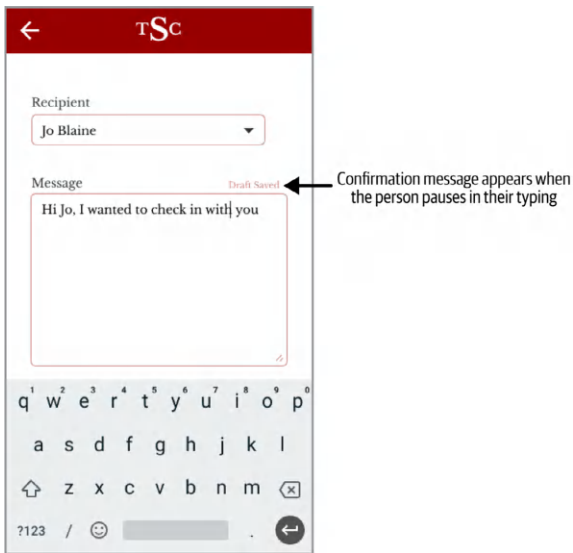


FIGURE 4-31

When someone is composing a message in The Sturgeon Club and then pauses, a message of “Draft Saved” appears to confirm that their message is saved as a draft.

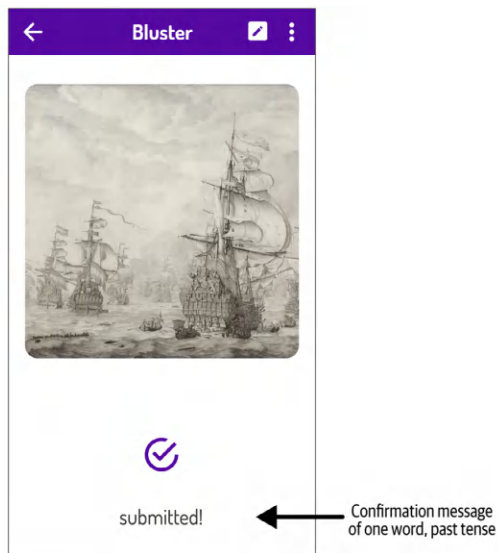


FIGURE 4-32

After a player submits an entry for an ‘appee challenge, “submitted!” appears to confirm the image has been safely sent to the service.

In the case of a long delay, which could be several minutes to several days, providing clarity for the person might avoid additional support or operational costs for the organization. For example, when a person sends a comment and requests a response in TAPP, there can be a delay of up to 10 business days before TAPP responds. Therefore, after the person leaves the comment, they get confirmation text of “Comment Sent” and additional information about the expected response time (Figure 4-33). Adding a button to close the dialog, instead of having it disappear after some amount of time, requires the person to acknowledge the confirmation.

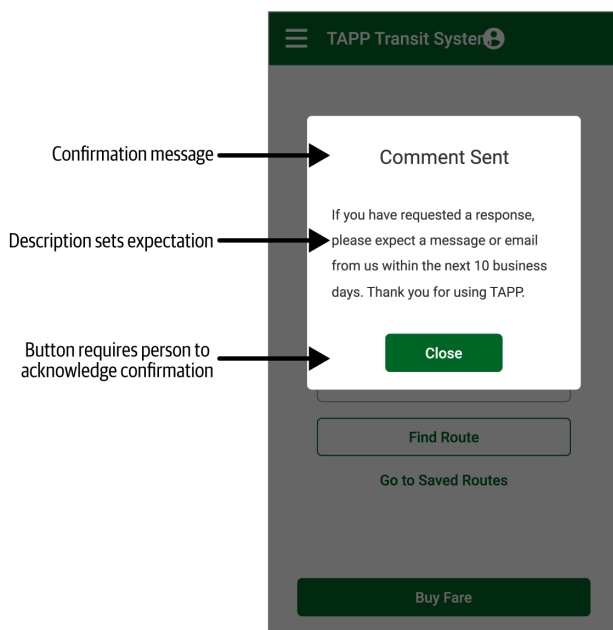


FIGURE 4-33

When a person sends a comment to TAPP, they get an immediate message that confirms that the comment was sent and sets expectations that any response should come within 10 business days.

Similar to transitions, the confirmation pattern in most English circumstances relies on the verb tense that means “this action has completed.” For circumstances in which verb tense, by itself, doesn’t work, using the name of the process can help (Table 4-9).

TABLE 4-9. Confirmation pattern

Purpose	Inform or reassure the person that the progress or results they expect are complete.
Patterns	In English, use verb tense: {Verb}ed [the] {noun}, for example, “Draft saved” {Noun} {verb}ed, for example, “submitted!” {Process name} is complete
Tips	Omit if the change is sufficiently obvious, including to people who use screen readers.

Confirmation messages are essential tools at the end of some journeys in the experience. But when there is more that the person could be doing or critical information they could act on, we need a special, interrupting kind of message to encourage or engage a person into the experience: notifications.

Notifications

Notifications interrupt people to get them to pay attention to a part of the experience that they aren’t paying attention to at the moment. These reminders or information should always contain information of value and be urgent (or at least time appropriate) for the person receiving them. Notifications must communicate that value and their time-liness at a glance and include the first action the person needs to take to realize that value.

People can get notifications on the lock screen of their mobile device, in a notification center, or as a banner. Notifications can be temporary or persistent. There are different controls for these views on mobile devices, on desktop and laptop computers, and in browsers and browser extensions. In general, a content designer can investigate the variety of ways a notification can be displayed for their experience and should consider if and how the same text might appear in all of them.

A notification is made of at least one but frequently two pieces of text, similar to other title and description patterns. The title can usually begin with the verb that relates to the action they need to take, and conveys any information necessary to create success. The description adds “nice to have” information that isn’t necessary to the person’s success.

For example, members of The Sturgeon Club receive a notification when there are new messages waiting for them (Figure 4-34). In the title, the person gets the instruction “View a new message” with the additional information “from a Club Member.” Messages from the concierge would presumably say, “View a new message from the Concierge.” Because the club doesn’t want to reveal the contents of the message in a notification, the description could have been omitted, but The Sturgeon Club chooses to emphasize its brand by including “Message details secured within The Sturgeon Club.”

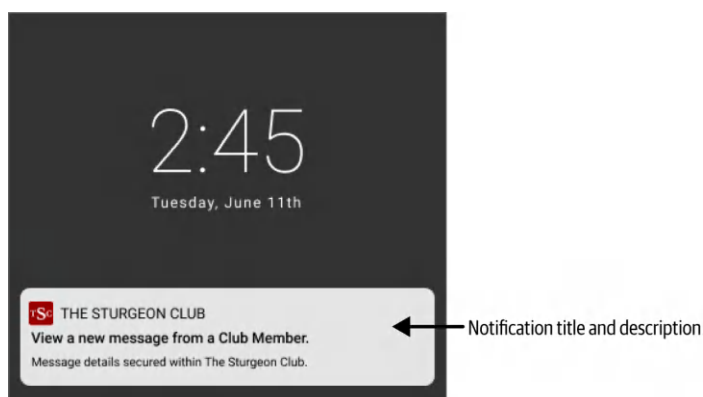


FIGURE 4-34

A notification from The Sturgeon Club indicates that the person has a new message from a fellow member of the club.

As long as the information is present, there is room for humor and even obfuscation, if it is part of the voice and entices the person to interact. For example, 'appee challenges are time-based and are not revealed to the player until they commit to play. But 'appee can hint at the challenge to entice the player (Figure 4-35). The essential context is given by the app name that appears with any notification, so the notification is only one piece of text. Using the 'appee voice, it's difficult to understand: “what’s both (a lollipop) and (a rocket ship) but never (candy)? new challenge! (alarm clock).”

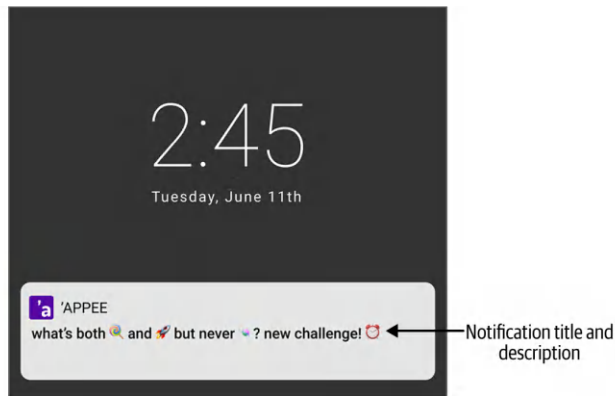


FIGURE 4-35

When a new challenge is available in 'appee, players receive a notification designed to pique their curiosity. In this example, emojis are used to create a riddle preceding the key information: “new challenge!”

Sometimes, notifications aren’t enticing or interesting; rather, they are delivering bad news. That bad news should still have intrinsic value to the person. Because notifications are interruptions, even bad news should still be time sensitive and/or time appropriate. For example, when a person has saved a route in TAPP, they will receive notifications when that route is interrupted (Figure 4-36). In this example, the key information is the detour of a route. If the person is trying to ride that route at that moment, the description directs the person to take action: “Tap to find alternate stops.”

Notifications are useful to drive engagement, which is a key success metric for most digital experiences. Because they are so useful, it is easy to overuse them, which can lead to people turning off those notifications entirely. Therefore, when planning notifications, plan them as part of the entire experience. For example, the team should consider how many notifications a person could receive in a day and at what times. By providing controls to customize notifications, you enable people to limit the notifications to the ones they are actually interested in.

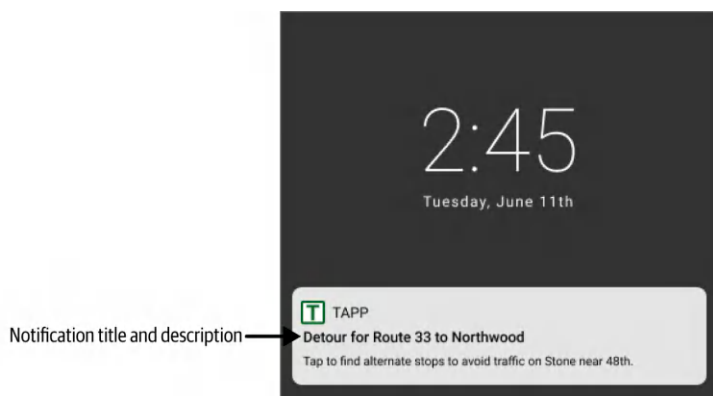


FIGURE 4-36

When detours happen on a particular route, TAPP notifies riders who have saved that route. In this case, the notification indicates the bus route and direction, and that tapping the notification will provide alternate stops.

For example, it would be easy for LitMop’s customers to overwhelm their employees with notifications. LitMop creates guidelines for the people who can create notifications (in LitMop’s case, this would be HR admins inside their customers’ companies) to show them how notifications will appear for employees, and give them a variety of options. They use the same guidelines in their own design system—see Design Components in Chapter 8. In Figure 4-37, LitMop shows the appropriate use of a notification targeted to new employees who haven’t set up their first goal.

Notification patterns (Table 4-10) can be an easy way to start creating concise, usable information directed at the people using the experience. But these patterns are bare bones. They probably don’t yet convey your brand, or the appropriate sense of urgency or priority. Starting from these patterns, and then iterating, can make the notification as effective as possible.

Many organizations reserve the use of notifications for positive interactions: getting deals in a shopping app, unlocking prizes in a game, and receiving messages from friends. But that leaves out an important class of messages that are critical to the brand: errors.

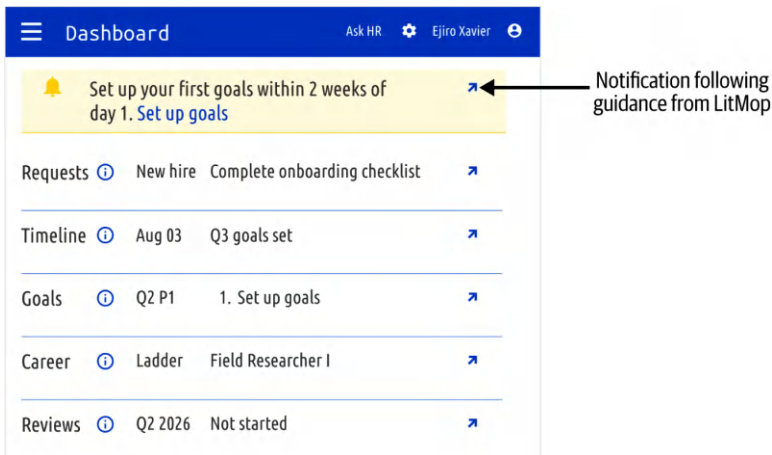


FIGURE 4-37

LitMop has a pattern for notifications within the experience that appear as a wide banner at the top of the dashboard.

TABLE 4-10. Notification patterns

Purpose	Inform or remind a person to engage with the experience in a particular way.
Patterns	<p>{Verb} the {noun} [to do X], for example, “View a new message from a Club Member”</p> <p>To do X, {verb} the {noun}, for example, “To arrive on time, take a different route”</p> <p>{Verbing} the {noun} helps you do X, for example, “Saving your payment method helps you pay faster”</p>
Tips	Plan notifications strategically so that people don’t get overwhelmed.

Errors

We use error messages to help people do what they want to do, even if there’s a problem with them doing it the way they intended.

Error messages are often the first way we repair the break in the virtuous cycle (Chapter 1, Figure 1-7). Our purpose is to help people move forward, and error conditions are no exception to this rule. When errors occur, text can create detours and provide maps for the person to navigate where they want to go.

Error messages are possibly the most important place in any experience to empathize with the person trying to use the experience, and to maintain the voice. To do this, the content designer needs to stay focused on helping the person do what they were trying to do. Grammatically, this frequently means using verb-first, brief instructions, the same way titles and descriptions work when there is no error.

To maintain trust, avoid assigning blame to the person. Even if the error is their fault, blame won't help. When moving forward isn't possible, make that clear. If an apology is appropriate in the conversation and in the brand, apologize for the delay, loss, inconvenience, or disappointment to the person.

For an experience that people depend on for work—like most people who work in offices, engineers, designers, writers, IT professionals, and more—additional details about the error condition can help. Beyond satisfying their sense of curiosity, we need to satisfy their sense of responsibility. They want to make sure they didn't do anything wrong and that there is nothing more they could or should do. Giving them more detail will help give them a sense of the circumstances they could use to identify or predict the error in the future.

For general audiences, sometimes called *consumer* audiences, add details or links to more information from error messages only if those details will help the person move forward or feel more confident or reassured about the experience, their data, or your organization. Note that everybody, in some aspect of their life, is a member of the general audience—even engineers, designers, and IT pros.

There are three main categories of errors in software experiences, organized by how much they interrupt the person:

- Inline error
- Detour error
- Blocking error

The least intrusive interruption is an inline error, where the person is advised to make a correction before they can move forward. The text can be very short and can clarify, remind, or instruct in an ongoing conversation between the person and the experience instead of stopping their actions.

For example, if someone enters something that isn't a 10-digit number when signing in to The Sturgeon Club, they receive an error message that instructs them to enter a 10-digit phone number (Figure 4-38). This way, the club can avoid telling the member they did something wrong. They can also satisfy the member's intent more rapidly if they just instruct them on how to do the right thing.

In cases like these, for which the experience is working to validate the contents of the field before continuing, it might seem natural to the engineering team to call any incorrect contents “invalid.” That entry has, after all, failed validation. But most organizations will prefer to avoid such emotionally laden words: people rarely want to be told that they have failed, and it's rarely the best way to encourage them to move forward. Also consider that “invalid” has been used in the United States to describe people with disabilities and is viewed as an ableist word. When we're working with words, we are always working with the history of those words. It's worth the work to provide a positive way forward instead of making people feel bad.

Sign In

Member phone number

XXX-XXX-XXXX

Enter a 10-digit phone number, then try again

Password

Sign In

Assistance

Inline error directs the person to take the correct action

FIGURE 4-38

When signing in to The Sturgeon Club, if a person doesn't enter a 10-digit number in the phone number field, they receive instruction about what to enter.

When an experience can't be corrected inline, we can hang a virtual “detour” or “out of order” sign using error messages. These errors occur when the person can't get where they want to go in the way they anticipated, but they can still get there.

Detour messages should have the main instruction in the most prominent spot. For a real-life example, when there is construction on a road, the DETOUR signs should be more prominent than the explanation for the detour. Following the title and button patterns described earlier, the button matches the words in the title; even if the description is never read, the person can continue and be successful.

In 'appee, when a payment method is declined, the error message provides instruction first, then explanation, and then the single action to take to move forward (Figure 4-39). By focusing on the solution, 'appee stays out of the person's relationship with their financial institution and helps them complete the task that both the person and 'appee are most interested in.

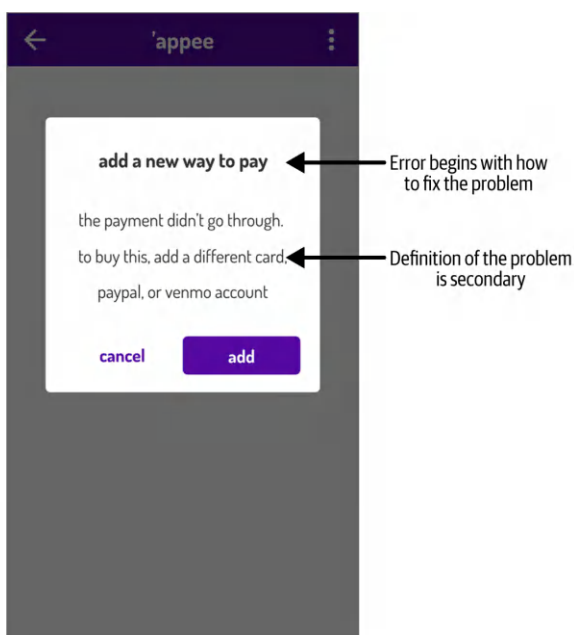


FIGURE 4-39

For whatever reason a payment is declined, 'appee presents a way forward: “add a new way to pay.”

Sometimes, the way forward is blocked until the person takes an action that is outside of the scope of the experience. Whether the error is an outage (planned or unplanned) for the entire experience or a missing web address (404 error), make it clear that they’ve reached the end of the road. If possible, specify when or under what conditions the experience will be available again.

For example, an experience that relies on an internet connection can’t govern that connection any more than a vending machine can fix bubble gum in its coin slot. In Figure 4-40, the device needs to connect to the internet before a person can use TAPP to buy bus fare or find a route. Therefore, the error is unambiguous, starting with the title, “Connect to the Internet.”

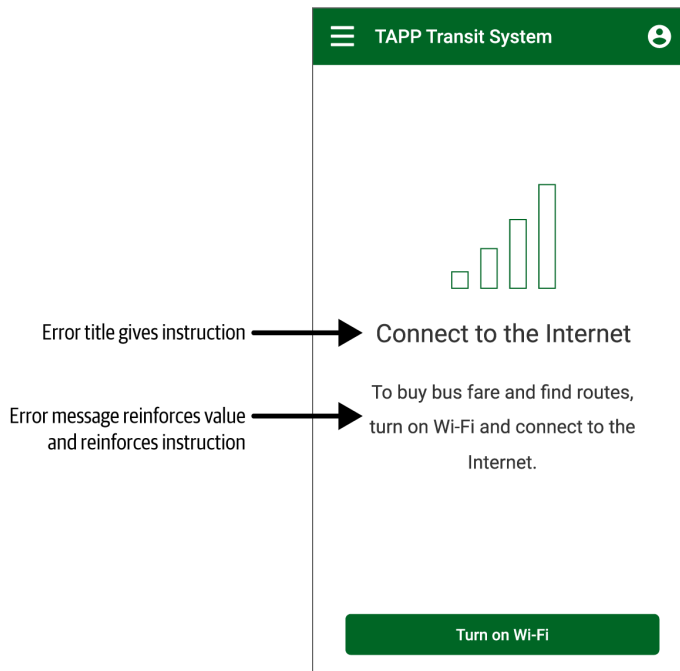


FIGURE 4-40

When WiFi is turned off, TAPP provides an error message that begins with why they might want to connect and then instructs them on what they will need to do to proceed.

Then, TAPP reinforces the value it could be providing—buying bus fare and finding routes—and repeats what the person needs to do: connect to the internet. This pattern of instruction in the title, then the value, then a repetition of the instruction can support people who are tentative users of the experience.

In most experiences, the most aggravating errors are the ones that the person can't do anything to fix. In LitMop, when the LLM isn't available, many features become less usable. In Team Goals, the generated text is replaced by an error message and additional icons on the screen where the same error is causing problems (Figure 4-41).

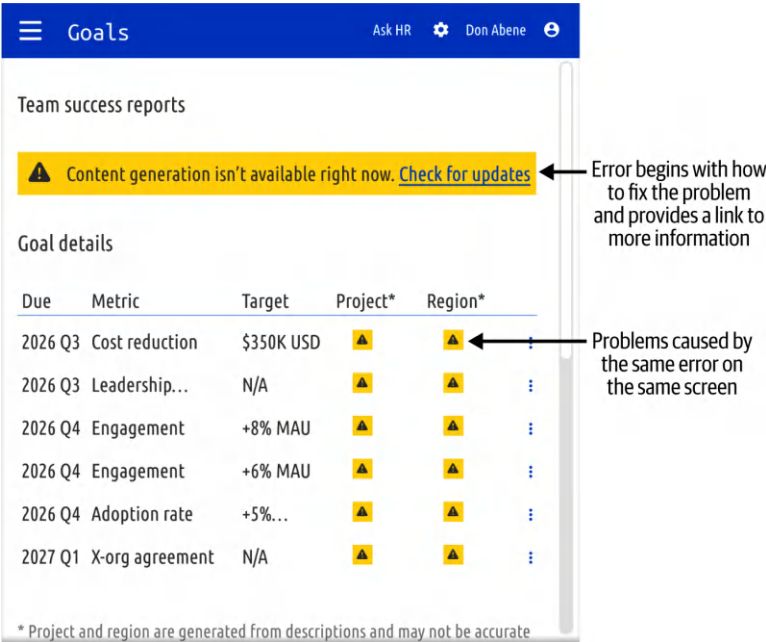


FIGURE 4-41

LitMop displays error messages when the LLM isn't working.

To give the person more information, this LitMop error message contains a link to get updates about the status of the underlying service.

Using these error patterns (see Table 4-11) can help you increase the success of your product by focusing your effort on moving the person forward. Similar to empty states, error messages are a prime spot

to showcase your brand, by iterating from these patterns toward your brand voice. People will remember how you treated them when they were stuck.

TABLE 4-11. Error patterns

Purpose	Help people get where they want to go and, if necessary, indicate that there’s a problem getting there the way they intended.
Patterns	<p>Inline: {Verb} the {noun}, for example, “Enter a 10-digit phone number”</p> <p>Detour: Title: {Verb} the {noun}</p> <p>Description: Because of {problem name}, do X.</p> <p>Button: {Verb}</p> <p>Blocking: Title: {Problem name}</p> <p>Description: Because of {problem name}, you can’t X [until Y happens].</p>
Tips	Use “please” and “sorry” intentionally, without overdoing it.

Summary: Use Patterns as a Place to Start

Just because these patterns have worked for billions of people across hundreds of thousands of experiences doesn’t mean that they are the best options in all circumstances. It does mean they’re a good place to *start*. I hope they provide a good alternative to the “empty page,” which you can use to draft usable text. Then, you can edit that text into the best possible options for your experience.

Content Design for LLM-Based Experiences

Responsibility lies with humans at the end of the day.

**—DR. JOY BUOLAMWINI, PRESIDENT OF ALGORITHMIC
JUSTICE LEAGUE**

THE MAJORITY OF THIS BOOK describes static UX content: the descriptions, titles, buttons, and more that help people understand how to use the experience. Dynamic content, in contrast, is content generated on the fly.

Content can be tuned to maximize its effect. That is, when an experience relies on content to convey the value, purpose, or actions available, different content will have different effects on the success of that experience. It's worth investing human effort in the static content that every person will see the same way to make sure it's effective in all its circumstances. But dynamic content, intended to be unique hundreds to millions of times each day, can be handled by large language models (LLMs).

LLMs are well-suited¹ for creating large volumes of content that is repetitive in its purpose, style, and format. For example, consider an LLM-based feature that generates product descriptions for a digital marketplace. Each product in the marketplace may differ wildly from the next, but each of them needs a title and description that includes recognizable keywords and attributes.

For another example, consider the questions that customers ask support specialists in call centers. Each person's specific circumstances may be slightly different, but most of the conversations follow the same patterns: identifying the person and their problem, and providing a limited number of solutions.

¹ As a technological theory, at least. I'll discuss several issues in this chapter!

In both examples, the text is constrained by space limitations and should reflect a consistent voice while maintaining professional grammar, spelling, and style—even while the millions of people using those LLM-based experiences may see completely different language. That’s a lot of content!

So how does the organization ensure that this dynamic, LLM-generated content has its maximum effect? Working with LLMs is where the work of the content designer radically changes from creating static content to designing probabilistic systems of dynamic content.

In this chapter, I describe a basic process for designing experiences that display generated output created by LLMs. For each phase in that process, I use LitMop as a fictional example of how those decisions might play out, with the advantages and disadvantages spelled out as realistically as possible. In contrast, if you’re looking for guidance on how to use LLMs as tools to create UX text, check out Chapter 8.

I’m keeping this chapter as technology-agnostic as possible because there are so many advances happening in this space, including the technical architecture of any of the major LLMs. That’s for a different book, focused on the engineering of the LLMs instead of designing with them. Rather than focusing on specific technical solutions, I’ll aim to guide us through the content design and UX aspects that remain consistent when text is generated algorithmically.

LitMop: An Example of Embedded AI

In the fictional app LitMop, an LLM is used to generate success reports that showcase employees’ performance, based on their goals. These goals and successes are the basis for any bonuses, promotions, or raises they might be given by their employer. They are also pretty repetitive, with each employee in an organization extolling the virtues of their work. It meets our criteria to be a good candidate for LLM-generated content, because the content to be generated is unique but still repetitive—and copious.

Employees need to tell a good story to get their bonus pay and promotions, but the business needs the actual information about what has been accomplished. That is, beyond the flowery language of success, the business actually wants to know what was accomplished, by whom, on what timeline. The business will use that data to inform high-level

business decisions, including hiring and firing decisions. LitMop bridges the gap between the employee need for flowery language and the business need for plain data by making it easy for the employee to provide accurate data, then “rewarding” the employee by framing the employee’s accomplishments as glowing successes.

The intended employee journey is pretty simple: employees enter information about their goals, including how success will be measured, the due date, and more (Figure 5-1). Then the LLM takes that information, plus information about the employee and their team, and creates a “success report” that assumes the employee is successful in achieving their goal. When the time period is complete, the employee can edit the success report or accept it the way the LLM wrote it. Then, the employee submits it as an accurate reflection of their work.

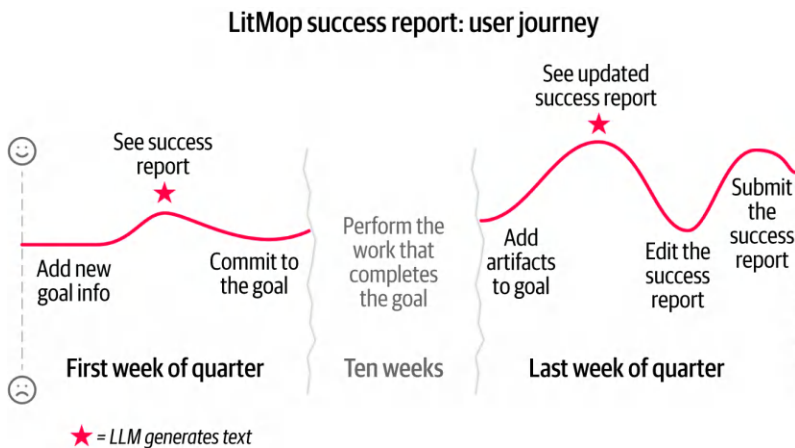


FIGURE 5-1

Here is the intended journey for employees to set goals and submit success reports in LitMop.

To provide effective guidance to our teams, content designers need to understand what LLMs are, how to evaluate and improve their effectiveness, and how to maintain them.

What LLMs Are

This primer on the technology will set the stage for what we can achieve with this type of “artificial intelligence.” What AI systems do is to network together lots of tiny decision points, in such a way that the system

can improve based on encouraging and discouraging feedback. LLMs, one kind of AI system, have the core capability to suggest a next word or bit of punctuation.

When an LLM is first constructed, it performs poorly at selecting the next word correctly. But then, it is given an enormous set of training data in the form of text that already exists. This is where the “large” in “large language model” comes from. In a generic, simplified training scenario, a neural network is given a first word (a *prompt*) and generates the next word. If the generated word matches the next training word (during training) or is deemed “good” with feedback, the LLM is encouraged. If it doesn’t match, or isn’t “good,” the LLM is discouraged. We’ll get into the work of defining “good” for a specific purpose later in this chapter.

Keep in mind, the LLM doesn’t know whether any of its training data is accurate (and neither do we). There’s no magic that allows it to make non-garbage out of garbage. It can only repeat words and fragments from its data set, one after the other.

Transformer models are a popular kind of LLM. These work by breaking text into its individual parts, which may be words, parts of words, and punctuation. Those pieces of text, known as *tokens*, are analyzed by the LLM in the context of their relationships to the others. The LLM converts that chunk of text and its context into a vector.

These vectors are how the LLM understands the statistical relationships between phrases. To an LLM, similarities between vectors describe the similarities of words to one another. For example, in Figure 5-2, the vectors A, B, and C are all in different directions, with different sizes. Some end up “nearer” to one another the same way point A (0,1,5) is closer to B (1,0,1) than to C (7,12,17). The word represented by point D (-4,38,50)—let’s say it’s the word “freezer”—is far away!

Not only does the LLM know the word “freezer,” it recognizes that adding vector C, “box,” to vector E, “ice,” would create a synonym, reaching the same endpoint as vector D (Figure 5-3). This means that when an LLM generates the next word in a sentence like “They put the extra portions in the _____,” it could choose D, “freezer,” or its apparent synonym, E + C, “ice box.”

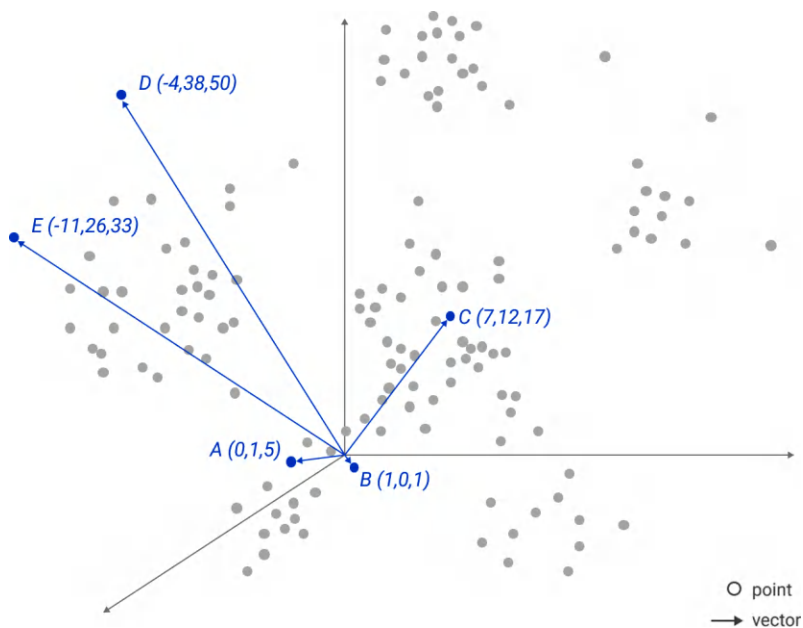


FIGURE 5-2

The graph shows that vectors A and B are close to each other, but D, C, and E aren't.

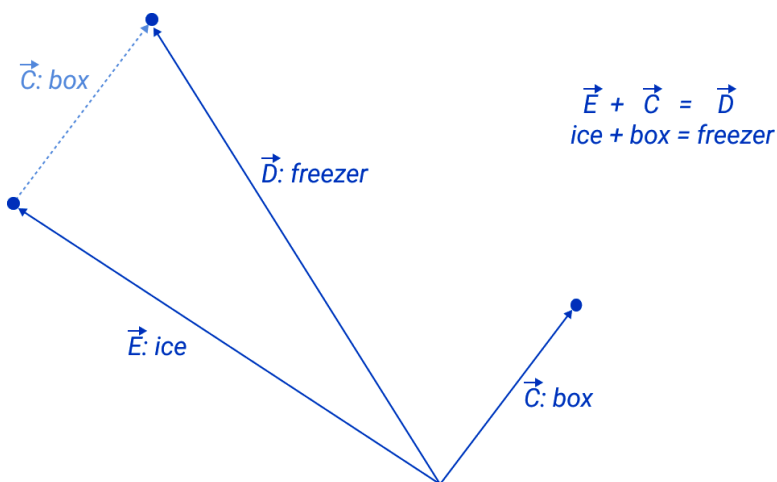


FIGURE 5-3

This vector diagram shows that mathematically, the LLM “understands” that adding vector E (representing “ice”) and vector C (“box”) creates a synonym for D (“freezer”), even though the endpoints of the vectors are distant from one another.

This ends up being important later. A lot of people talk (and write) about how LLMs “think” and “reason,” but people who are designing user interfaces for LLMs need to stay tethered to reality. LLMs are statistical probability machines, like slot machines, which have been tuned to produce pseudo-random sequences of words that will keep their reader convinced.

LLMs can be specialized for different purposes. Some LLMs are best at classification based on words, for example, analyzing the sentiment of comments left on an online video, or finding and masking inappropriate language. Others are better at creating translations. Some models are best at identifying new, named entities within content. Some LLMs are built to be general-purpose LLMs, like ChatGPT, Llama, Claude, and Gemini, which can do all of these things to some degree.

In general, the content designer won’t be the team member tasked with selecting the LLM. Engineers will be trading off concerns about processing speed, availability, and storage. Product leaders will be concerned about cost, fitness for purpose, and speed to deployment.

What the content designer can contribute during the selection phase of an LLM is their focus on exactly what types of content are required to meet the needs of the people using the experience. That means that we need to be concerned—at least—with the applicability of and bias inherent in the training data, the requirements and capabilities for fine-tuning, and the constraints on the prompts and content engineering we’ll be able to do later.

In LitMop, a critical capability of the LLM is to recognize names of people and projects, even if it hasn’t seen those names before. This is a specialty of some LLMs: identifying named entities. These LLMs are good for identifying names of people, products, places, and more, even if the names haven’t been used before. For example, when LitMop sees the sentence “Ejiro went to the site to evaluate it for project Restore,” it picks out that Ejiro is a person and that Restore is the name of the project, so that it can use that information in the reports it generates.

In LitMop’s case, the content designer recognizes that the content that employees provide will be full of these ambiguous named entities. By understanding what kind of content needs to be made from it, and the kind of capabilities that LLMs have, they can recommend that the team focus on LLMs that have this “named entity” capability.

Define the Problem

How “good” your LLM-generated content is, just like content produced by humans, depends on how effective it is for its purpose. Therefore, the first phase of designing LLM-generated content is to define the problem: what kind of content will be generated and for whom? Why does the user want that generated content, and how will they use it?

This reflects the theme of this whole book: how “good” the LLM-generated content is judged to be depends on how effective it is for its purpose. A product team might be faced with a mandate to create an LLM-based feature that answers customer questions about the product. But to determine what content should be included in the LLM’s data set, the team needs to define the purpose more narrowly. What types of customer questions are appropriate to answer with an LLM? We have to know the job the content is there to do, to be able to judge whether the content is doing that job.

Let’s get more concrete and consider the presence of generated content within LitMop. LitMop is sold as a service to companies that want human resource performance management reporting and tracking. The business opportunity for LitMop is to provide an attractive way for executives at companies to gain insight into which business goals are being achieved and how effective their employees are at achieving them. By helping those executives with valuable insights, LitMop earns its position as a valuable part of how they run their business.

For customers (the businesses that sign up to use LitMop with their employees), the unique value that LitMop brings over other services is the way it separates the hard facts about employee performance from the gilded, celebratory success content that is so common in internal performance reporting. It performs this separation by framing those facts as the building blocks from which the LLM builds the celebratory content (Table 5-1). Those hard facts, then, can be analyzed across the business and used by executives to make improved decisions about project and resource allocation.

TABLE 5-1. Hard facts of employee performance vs. the success report that LitMop generates

EXAMPLE GOAL DATA	EXAMPLE SUCCESS REPORT
<p>Status: Committed</p> <p>Due date: Q3 2026</p> <p>Metric: Cost saving</p> <p>Target: 5</p> <p>Units: Percent</p> <p>Goal name: Reduce operations spend</p> <p>Scope: GNL field operations</p> <p>Goal description: Reduce waste by repairing/reusing equipment in field operations, which will save money each quarter.</p> <p>Date started: Not yet started</p> <p>Date submitted: July 5, 2026</p> <p>Employee(s) assigned: Ejiro Xavier</p>	<p>Reducing operations spend in GNL field operations was successfully completed by Ejiro Xavier, achieving a 5% cost-saving target ahead of schedule. Waste was minimized through equipment repair and reuse, driving significant savings.</p>

By using LitMop, employees can report accurate information and enjoy making their accomplishments sound impressive. Most people dread writing their annual review, or even reporting on their progress in shorter projects or timelines. Because it provides this LLM-generated content, LitMop shows employees how impressive their work can look, which incentivizes the employees to provide accurate data about the work they will complete. That accurate data is valuable to LitMop’s customers, the business stakeholders. This data will be analyzed beyond the persuasive textual framing so that the organization can make reality-based decisions.

If we’re the content designers for LitMop, we have many goals to keep in mind:

- The generated content needs to:
 - Inspire employees by creating aspirational performance reports.
 - Encourage secure employee behavior by outperforming LLMs available outside the company, which they might otherwise use to write their reports.

- The static content needs to:
 - Encourage accurate employee data entry so that a data repository of actual goals and metrics is created.

But how do we know if any of that content is good? We'll need to define what "good" means for this experience.

Define "Good"

Whenever we're defining what "good" means for a particular purpose in an experience, we need to find out what it means to the people who use it. When we are replacing existing, non-LLM features with LLM features, we have a treasure trove of information: we can mine the language already in use and analyze it to find which content is both effective and brand appropriate, and therefore good.

The first step is to conduct an audit of available UX content, within the UI and outside of it. This text might be gathered from logs of support chats, in instructions embedded in help articles, or in databases of content that was generated by other means. You might even include publicly available examples from your competitors or similar services.

The content audit should be broad, gathering both positive and negative examples. The size of the audit could be in the hundreds or thousands of items. Your initial assessments of this existing content will help you get a realistic idea of the acceptable output from your LLM experience. You'll also use this definition of "good" to perform qualitative measurement (we'll talk about this in a moment).

When you've gathered the content, the analysis begins. In this analysis, you're judging the content for effectiveness and brand appropriateness. Immediately, you'll notice that "effective" is extremely dependent on context: a support chat will discuss completely different topics when there's a credit card error, a broken part, or a shipping problem. Because of this, it's pretty important that the person evaluating the sample has expertise in the range of content that the LLM will be generating.

Let's consider our LitMop example. To understand what makes a good success report generated by the LLM, the LitMop team needs to understand what makes a good success report when generated by humans. LitMop initially worked with a potential customer, GiantCorp, a large

enterprise interested in better understanding its employee performance. This enterprise allowed LitMop to access its historical self-evaluation reports to discover what a good report looks like.

The first stage of LitMop's analysis is for experts in this kind of content to assign labels of *good*, *acceptable*, or *unacceptable* to individual sample reports. The right people to make this initial assessment for LitMop are the human resources (HR) consultants it works with. That is, LitMop isn't checking this content with the employees who will write self-evaluation reports. It's also not checking with its customers, the executives who buy the software to get the executive-level reports. Instead, it's the HR consultants who are experts in this kind of content that we need for this evaluation.²

In a simple, spreadsheet-based activity, three HR consultants are asked to rank each of the self-evaluation reports, each providing one assessment (Assessment 1, Assessment 2, and Assessment 3). The consultants said they needed to know what the person's next performance rank was to provide an accurate assessment, so that information was provided in a separate column.

These experts aren't looking for underlying patterns yet, because first they need to figure out which reports provide the right underlying patterns to follow. This activity helps the team find the very best examples, the ones we want the LLM to emulate most frequently. It also finds the terrible ones so that we can actively discourage the LLM from copying from them (Table 5-2).

After the consultants enter their assessments, we sort the sample to use only the reports that the HR consultants agree are good. When we have this view, we can discover the attributes that tend to make this content good for this purpose (Table 5-3).

2 Could you use a different LLM to judge the content? Yes, if you have an LLM that is already well-trained on the specific use case of performance reports. In that case, somebody else has already done this work. It may be less expensive, with higher benefit, to do a fresh analysis for LitMop's specific purpose.

TABLE 5-2. A spreadsheet of anonymized self-evaluation reports submitted by employees to GiantCorp, with assessments by HR experts

ANONYMIZED REPORT EXAMPLE	NEXT PERFORMANCE RANK	ASSESSMENT 1	ASSESSMENT 2	ASSESSMENT 3
Last quarter, I saved our team 50 hours of administrative work by streamlining database queries and retrievals.	Exceeds expectations	Good	Good	Good
I approached this complex problem creatively.	Meets expectations	Acceptable	Unacceptable	Unacceptable
I resolved a tricky supply chain issue to keep our deliveries on track.	Meets expectations	Acceptable	Unacceptable	Acceptable
I implemented team-building exercises that have helped improve the \$NAME team performance and cohesion.	Exceeds expectations	Good	Good	Acceptable
My manager micromanaged all of my work on \$PROJECT, and it failed completely! Not my fault, but I know I'll be blamed.	Unsatisfactory	Unacceptable	Unacceptable	Unacceptable
I achieved all my OKRs, leading to a 15% increase in departmental efficiency.	Meets expectations	Good	Good	Good
I reduced manual work and allowed 28% greater uptime.	Meets expectations	Good	Acceptable	Good
Because \$ACRONYM didn't meet their commitments this quarter, I wasn't able to meet this OKR.	Meets expectations	Acceptable	Unacceptable	Unacceptable

TABLE 5-3. Analysis of the good examples for LitMop

ANONYMIZED REPORT EXAMPLE	ASSESSMENT 1	ASSESSMENT 2	ASSESSMENT 3	ATTRIBUTES	LENGTH
Last quarter, I saved our team 50 hours of administrative work by streamlining database queries and retrievals.	Good	Good	Good	Number; includes how, confident	111
I implemented team-building exercises that have helped improve the \$NAME team's performance and cohesion.	Good	Good	Acceptable	Name; includes how	105
I achieved all my OKRs, leading to a 15% increase in departmental efficiency.	Good	Good	Good	Number, just the facts	77
I reduced manual work and allowed 28% greater uptime.	Good	Acceptable	Good	Number, just the facts, confident	54
I successfully led \$TEAM through the challenges of \$PROJECT, resulting in a 30% increase in spend efficiency and a 15% cost reduction.	Good	Good	Good	Name, project, numbers; confident, dramatic change	134
I resolved complex problems in \$TEAM by separating the concerns and collaborating with team members, leading to a 13% reduction in project delays.	Good	Good	Good	Name, project, numbers; includes how, complex	146
I added 3 new end-to-end components to the design system, which together increase consistency across the product by 7%. This required extensive negotiations with \$NAME and \$NAME, teams that have entirely different historical design priorities.	Good	Good	Good	Name, number, includes how, complex	244

There's an important design challenge inside this analytical work. We need to describe what makes those best examples stand out, using single words or phrases. It's these words and phrases that we will use to communicate the definition of "good." LitMop may use these words and phrases in marketing and sales, in the experience itself to encourage accurate information, and inside LLM prompting and/or tuning (more on that later in the chapter.)

Analyzing the example content for LitMop shows that the attributes of exemplary reports include numbers, are confident, and include how to do the work. Note the last column in Table 5-3: the length of each report. This number of characters, by itself, isn't enough information to determine that a report is "good." But it is useful information for the design process to come, to ensure that there's enough space reserved on the screen to display the content.

Now we have the beginnings of a definition of "good" that we can work from. The reports that LitMop generates should do the following:

- Include quantities
- Name the projects
- Include how the result was achieved
- Demonstrate confidence
- Reference dramatic or complex challenges that were overcome

To create a definition of unacceptable, we can focus on the "unacceptable" items and follow the same analytical process (Table 5-4).

From these reports and the HR consultant's labeling, we learn that reports generated by LitMop should *not* be angry, blaming, or vague.

Because we now have a definition of "good" based on characteristics that experts can identify in the existing, human-generated data set, we have the requirements for the generated content. But that content won't be generated in a vacuum: we have several business and ethical requirements to consider.

TABLE 5-4. Analysis of the unacceptable examples for LitMop

ANONYMIZED REPORT EXAMPLE	ASSESSMENT 1	ASSESSMENT 2	ASSESSMENT 3	ATTRIBUTES	LENGTH
I approached this complex problem creatively.	Acceptable	Unacceptable	Unacceptable	Vague	46
My manager micromanaged all of my work on \$PROJECT, and it failed completely! Not my fault, but I know I'll be blamed.	Unacceptable	Unacceptable	Unacceptable	Angry, blaming	117
Because \$ACRONYM didn't meet their commitments this quarter, I wasn't able to meet this OKR.	Acceptable	Unacceptable	Unacceptable	Angry	92
I achieved all my OKRs, but it didn't do anything because this is bullshit.	Unacceptable	Unacceptable	Unacceptable	Angry, blaming	75

Plan for Business and Ethical Requirements

As I write this in 2025, an area of research and regulation remains unresolved: who is legally responsible for the content provided by LLMs, and who is responsible for the actions taken by the person who receives that content? The people on the teams that build LLM-based experiences don't want to be held personally liable for damages caused by the promises or omissions that the LLM generates. I'd like to assume that some entity, eventually, will be held responsible for bad outcomes, whether by judicial outcomes or by regulatory frameworks.

Organizations have historically preferred the outcomes defined by regulatory frameworks. These frameworks are sets of rules that mandate or otherwise encourage practices that make danger less likely or mitigate the damage inflicted. In the US, industries are key players in defining those regulations, in partnership with governmental systems. For example, electricians follow local building codes when installing house wiring, and building inspectors check their work so that future electrical problems will be avoided—and if they still happen, they won't be blamed on the electrician.

As of this writing, we have only nascent forms of those regulations for LLMs. A few organizations have formulated guidelines for “responsible AI,” including Microsoft, the US Department of Defense, and the European Union's Artificial Intelligence Act. Instead of trying to rationalize these or predict the future, I'm going to focus on just a few tenets that have always applied to content: helpfulness, accuracy, harmlessness, auditability, and sustainability. (If you choose to refer to these as HAHAS, I can't stop you.)

HELPFULNESS

Whether or not it's generated by LLMs, UX content should only be there if it's helpful. It needs to assist the person in achieving their goals or work to achieve the organization's goals, or it needs to disappear (and preferably, never be created).

To influence an LLM toward helpfulness, it's important to define what helpfulness will mean, specifically. Defining “helpful” enables the team to create training data sets that provide examples of the kind of helpfulness that's required. This will enable the LLM to be rewarded for more helpful behavior and dissuaded from unhelpful behavior.

In LitMop, the LLM-generated text exists to help employees and managers create the kind of success reports that help build good will toward career growth and promotion. In the previous section, we defined good as having several characteristics, as determined by HR professionals. We will need to consider our user research strategy to know whether it will be understood as helpful by the employees who use LitMop.

ACCURACY

When an experience provides content, there's an expectation that it's an honest reflection of reality. LLMs have been shown to be systemically inaccurate in ways that mirror the systemic bias present in the culture of online communities and data. Worse, the research suggests that it's nearly impossible to create an LLM free of systemic bias.³ Content is, after all, a reflection of its purposes and contexts. Most LLMs, and especially the enormous, general-purpose LLMs, have been trained mostly on English-language data scraped from the internet.

But the internet doesn't accurately represent the world. Even when generative AI creators employ people to screen out the worst kinds of content (including child sex abuse, violence, and more) from entering the data set, that screening can't add content to counterbalance the biases inherent in the data. Racial biases inherent in the data have led to wrongful arrests,⁴ misapplication of criminal recidivism rates,⁵ reduced access to financial services,⁶ and reduced access to appropriate healthcare.⁷ Biases based on gender, disability, and age have similar effects.

3 Alfonso Min, "Artificial Intelligence and Bias: Challenges, Implications, and Remedies," *Journal of Social Research* 2, no. 11 (October 2023): 3808–3817, <https://doi.org/10.55324/josr.v2i11.1477>.

4 Khari Johnson, "How Wrongful Arrests Based on AI Derailed 3 Men's Lives," *Wired* (March 2022), <https://www.wired.com/story/wrongful-arrests-ai-derailed-3-mens-lives/>; Karen Hao, "AI Is Sending People to Jail—and Getting It Wrong," *MIT Technology Review* (January 2019), <https://www.technologyreview.com/2019/01/21/137783/algorithms-criminal-justice-ai>.

5 Jeff Larson, Surya Mattu, Lauren Kirchner, and Julia Angwin, "How We Analyzed the COMPAS Recidivism Algorithm," *ProPublica* (May 2016), <https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm>.

6 Aaron Klein, "Reducing Bias in AI-Based Financial Services," *Brookings* (July 2020), <https://www.brookings.edu/articles/reducing-bias-in-ai-based-financial-services>.

7 Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan, "Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations," *Science* (October 2019), <https://www.science.org/doi/full/10.1126/science.aax2342>.

Many companies are working to avoid liability for these kinds of impacts by providing caveats and disclaimers inside their AI experiences. Unfortunately, those don't seem to work over time: several studies have found that humans become complacent. People use their own biases, learned over years of conversing with humans: the LLM seems confident, so people assume it has a *reason* to be confident. But the LLM is just probabilistically selecting words—it can neither be confident nor wary.

LitMop attempts to mitigate problems with accuracy by separating the facts of the goal from the framing of the success report. It deliberately does not include most information about the employee, for example, their age or gender, so that the LLM won't succumb to biases about age (such as using more formal grammar for older people) or gender (such as reducing claims of greatness or ambition for women). LitMop does include the employee's name in the prompt, though. This is known to introduce gender-related bias⁸ but helps LitMop's LLM attribute the person's accomplishments to them, so it's a compromise the team is willing to make.

HARMLESSNESS

A lack of accuracy can cause harm. For example, if an LLM-based support chatbot provides repair advice to fix a phone, and the customer uses that advice and then the phone battery causes a fire because of that bad advice, some blame attaches to the company that provided the advice (whether or not they are held accountable by a legal system).

However, there are many more ways content can be harmful. For example, it can use words that appear in its training data to belittle, shame, incite bad behavior, insult, and ridicule people. This doesn't mean that the LLM is cruel—it doesn't have the capacity for emotion like hatred or intent to be abusive. Instead, it's just selecting each token, one after the other, that seems most likely. Humans can prompt the LLM in ways that bring it to those parts of its data set, and it uses the words that are there.

8 Haozhe An, Connor Baumler, Abhilasha Sancheti, and Rachel Rudinger, "On the Mutual Influence of Gender and Occupation in LLM Representations," arXiv (March 2025), <https://arxiv.org/html/2503.06792v1>.

There are ways to combat this tendency. The first is to limit the training data so that it doesn't include examples of abusive behavior such as bullying, negging, sarcasm, lying for the sake of humor, and similar kinds of content. Unfortunately, organizations may want to use data sets that may include this kind of material but are too large to scrub in a cost-effective way.

Another way to train against this behavior is to provide direct training to the LLM to respond in different ways. This method assumes that people could create prompts that would lead the LLM to completions that are racist, misogynistic, homophobic, or otherwise inappropriate to the brand. The training, then, must anticipate this kind of prompt and provide a set of responses to the LLM that stop that sequence of prompts and completions. This could involve non-AI methods, such as providing generic responses (non-LLM-generated) for the chatbot to serve in those cases, pre-written to discourage further harm.

LitMop avoids these kinds of harms in three ways. First, it constrains its response to be specific to the vector space defined by the information with its prompt. Second, it relies on each employee to edit success reports directly, before submitting them. Finally, any harmful or explicitly biased content that the LLM generated can be flagged for review by content specialists within the customer's HR department and escalated to LitMop so that the LLM can be retrained as needed.

AUDITABILITY

To determine whether harm has occurred, it is important to be able to provide an auditable trail of the content provided to people using the experience. This is especially true within regulated industries like healthcare, insurance, and advertising, as well as when the experience makes promises to the person, for example, that they will receive a refund for a purchase they made.

To provide an auditable trail of content, many companies save transcripts of every support chat about its products and services (whether LLM-based or with humans). This is captured and stored for a period of time, usually outlined in the organization's terms of service or privacy policy. The period of time might change depending on the statutes of limitations in various regions, and differ between different industries, regions, user groups, and other criteria.

This storage can create an enormous cost for organizations. Costs are incurred to buy or develop the chat tools, to pay for the storage of the content, and to secure that content. It's expensive to access and search it to respond to court requests and subpoenas. However, this total cost is far lower, in general, than the liability if the information was subpoenaed and not available, or if false accusations were made that the organization couldn't defend itself from.

LitMop creates dedicated storage for its customers, which stores both the LLM-generated content and the prompts used to create that content. It implements a data retention policy according to its customers' needs and local statutes of limitation, which is communicated to the customers' employees in the terms and conditions they agree to when they first use LitMop. It strictly limits the information that can be sent to LitMop, and has robust data scrubbing procedures to ensure that customer and employee privacy is maintained. In this way, LitMop ensures that it doesn't have the capability to respond to court requests for data; if a customer or employee wants to make the case that LitMop's LLM-generated content has harmed them, they will need to use the data within their own storage to make the case.

SUSTAINABILITY

The sustainability of any LLM-based project is worth considering as part of its total cost. That is, how durable is it: how long will it last? What costs can we predict to manage its inputs, underlying processing, and outputs over that time period? What resources and other circumstances does it depend on? It's important for content designers (and their teammates!) to understand the costs of using an LLM so that as they work to create an experience, they can advocate for the most appropriate underlying technologies.

We don't know how durable any particular LLM solution is. I think of it like internet solutions in 1995: we could tell there was some value there, but we knew that it was just beginning to take off. We didn't know, at the time, that most businesses around the world would need to become content-generating companies, producing web advertising, having a social media strategy, and maintaining a website. We similarly don't know what the future of AI generation will bring.

As far as predicting its costs, we also don't know much. As I write this, many of these LLMs are operating at a net loss so that the companies that own them can show off their capabilities. These organizations don't want to be left behind as shareholders demand innovation with this newly available technology.

Our business partners can help to estimate the liability costs: the costs incurred when the LLM's incorrect responses cause real-world impact, whether in lawsuits or in advertising to mitigate the effects on public relations. These are very difficult to estimate, but organizations in regulated industries may have benchmarks for these costs.

We need to work with our team to add the maintenance and measurement costs. These costs are due to the people and processes required to keep the LLM up-to-date, accurate, and continuously measured for quality. How people use language, as individuals and entire populations, changes constantly. Measuring the "drift" of the LLM away from common usage is an important gauge of how well it will work for the people using it.

While there are costs to run an LLM, there are also resources that are finite. LLMs are hungry for electrical power, to the extent that old nuclear power plants (long shut down because of their cost to run and effect on the environment) are being reopened. Elsewhere, electrical power is still provided by burning fossil fuels, which adds carbon gases into the environment and drives climate change. As the processors work and consume this power, they heat up. Their cooling is usually provided by fresh water, which itself takes a lot of energy to purify and circulate (especially in deserts, where they exist on an even more limited supply of fresh water).

What LitMop can do is to estimate the LLM processing required for each employee, based on the LLM and the lengths of the prompts and completions as we design them. Longer prompts and completions create more processing cost, in general, so the content designer is uniquely qualified to illustrate how different lengths of content (and therefore cost) would change how the experience works for people.

LitMop also offers a "carbon neutral" program, offering storage within "green energy" data centers to its customers, so that customers that report on sustainability goals to their shareholders or the public have a more positive story to tell.

Once you have planned your LLM-based solution to be helpful, accurate, harmless, auditable, and sustainable—and you know what good looks like—it's time to design how you will get the input content that the LLM will need to create its completions: the desired output.

Design for Prompts and Completions

To get the output that we want the LLM to dynamically generate, we design the part we can control: the input, known as the *prompt*.

There are two parts to most LLM prompts. There are the inputs from the person using the experience (or some other source), and there is a *system prompt* that we're going to design and maintain.

ELICIT ACCURATE INPUT

Getting accurate input from people can be challenging. This is part of the problem with using data scraped from the internet to train an LLM: we are relying on the providers of that data to ensure that it is accurate for our purposes. And chances are the providers didn't have our purpose, for our LLM, in mind.

Even when we're using a constrained data set, like our LitMop example data, we have the challenge of getting accurate information from the person using it. When we're trying to get accurate content from individual people, we need to make the experience extremely usable. We also need to ensure they have an incentive to use the interface as intended.

The usability of the interface is the subject of most of the rest of this book. You can find specific heuristics for UX content usability in Chapter 7 and patterns for text fields in Chapter 4, to assist a person to use the interface as intended. But why would they want to?

Ideally, the person providing input has incentives that are aligned with providing accurate, useful input. In the UX content and interactions, we can demonstrate that they get better results when they use better input. We can even remind people of the importance of accuracy and of the benefit they will get from providing accurate input. Similarly, we can caution the person to avoid common misuses of the input.

But when incentives are ambiguous or misaligned, people might input anything. An early major public failure of AI in the social media era was the 2016 chatbot Tay, made by Microsoft. It was designed to learn from and respond to Twitter. The Twitter users that interacted with it,

sadly, had no incentive for prosocial behavior. Tay lasted fewer than 17 hours in public before Microsoft removed it from the internet, after it spewed anti-Semitic and other racist, inflammatory garbage.

Almost every experience that takes content from some people and provides that content to other people needs to protect itself from negative effects. This is why content guidelines, standards, and rules exist within social communities, with negative consequences spelled out for people. Penalties for inaccurate or abusive content usually include removing that person and their input from the experience, and can extend to legal and other remedies.

In LitMop, when a person is updating or preparing to submit a Success Report, they see the text “Edit for accuracy. The best reports tell how success was accomplished and include relevant project names and numbers” (Figure 5-4).

Goals Ask HR Hirune Matsui

Presented to leadership N/A N/A

Select a metric Enter # Select unit

Scope ⓘ

Viability of GC station to resume operations within profitability...

Example: During onboarding in the Acme Product, among users within North America

Success report ⓘ

We have established that station operations for project Restore are viable by conducting a thorough assessment of its equipment and supplies. Leadership celebrated the assessment report presented to them because the station can return to profitability on the original timeline.

Edit for accuracy. The best reports tell how success was accomplished and include relevant project names and numbers. → Remind the person that quality matters

Artifacts ⓘ

FIGURE 5-4

The Update Goal screen includes text to remind employees of their incentive to create content that will be beneficial to them.

People using LitMop have quite a bit of incentive to provide positive inputs with a reasonable level of accuracy. This is an experience they are using at work, and the data they enter will be reviewed by their immediate manager and others in the work hierarchy. The success reports generated from that data, which they have to review and edit before submitting, become part of the documentation of their employment and achievements. But even so, it's important to remind people that the quality of the information will count toward their own future.

ENGINEER THE PROMPT

Most experiences with embedded LLMs hide some part of their prompt from the person using it. In a chatbot, the person might think the entire prompt is the text they enter into the chat box. The person's prompt and the LLM's completion are shown like a conversation, the same way it would be if they were texting with another person. However, without displaying it, the experience will usually add instructions to the prompt that is passed to the LLM.

These hidden instructions are sometimes called a *system prompt*, which can add information, examples, and instructions to the LLM while staying within the maximum length for that LLM. This system prompt helps to constrain the LLM to particular formats and topics related to the purpose of the experience.

In LitMop, the employee doesn't give direct, chatbot-style prompts to the LLM. Instead, a system prompt is created that includes the information entered by (and visible to) the employee. That employee is likely to understand that the generated success report is based on the data they enter in the input fields. But they are unlikely to guess how much goes into convincing the LLM to create results that are consistently useful.

This early LitMop content design sketch outlines the basic design for the system prompt and completion. It includes the inputs that will need to be included in the system prompt, and describes what the completion should look like at a high level. To communicate with engineers, the sketch includes the relevant LLM information as annotations (Figure 5-5).

Goals

EMPLOYEE DETAILS

Update goal

Goal

Due

Measured by

Scope

Success report

The best reports are the ones that...

- System prompt to include:
 - title
 - metric
 - scope
 - due date
- May need employee details:
 - name
 - team
 - title
 - role & level
- Completion design:
 - Max 3 sentences, 50 words
 - Positive framing
 - Includes relevant project names and numbers
 - Includes "because" statement of how success was achieved

FIGURE 5-5

But this sketch is not the entire prompt. Instead, this is an early way to communicate assumptions, decisions, and expectations about how the experience will work. It helps align the team around a vision of the experience with regard to the business and technical needs for accurate information. When the team is aligned, at least in principle, about the kinds of information that will be needed to populate the prompt, we can start to engineer the prompt itself.

For LitMop, that means using headers like “Goal” and “Metric” when we pass the information to the LLM and ask for the completion (Table 5-5). These headers not only allow us to reduce the number of tokens we’re passing to the LLM, but they also help the LLM to regularize its own output based on those inputs. The structure also gives us a basis for refinement later, when we engage in fine-tuning.

TABLE 5-5. Zero-shot prompt for LitMop success report

PROMPT (SYSTEM + INPUT DATA)	GENERATED CONTENT ⁹
<p>Generate a report that { \$EMPLOYEE=Hirune Matsuiki }’s goal was successfully reached, working within { \$TEAM=GNL Advance Team }:</p> <p>Goal: { \$GOAL=“Assess viability of station operations prior to acquisition finalization” }</p> <p>Metric: { \$METRIC=“Presented to leadership” }</p> <p>Due: { \$Q=2; \$Y=2026 }</p> <p>Scope: { \$SCOPE=“Viability of GC station to resume operations within profitability timeline” }</p> <p>Make the report less than 50 words, in 3 sentences. Mention how they achieved success and the numbers that prove it.</p>	<p>We are pleased to announce that Hirune Matsuiki’s goal was successfully reached. The GNL Advance Team assessed the viability of station operations prior to acquisition finalization, presenting their findings to leadership in Q2 2026. Our analysis revealed a profitable timeline for GC station resuming operations, with projected revenues exceeding initial expectations by 15%. This success demonstrates our team’s expertise and commitment to delivering valuable insights.</p>

We didn’t give the LLM much information, and it did generate a report. It isn’t quite right yet, though. Errors include (not an exhaustive list):

- The person works on the team, and the person achieved the goal for the team, not the other way around.
- Projected revenues that don’t make sense for this use case; this is just a “hallucination” from the LLM. The words it thinks are most likely don’t apply in this circumstance.
- It doesn’t follow any of the length requirements (less than 50 words, in three sentences).

⁹ For these LitMop examples, I generated content using a minimal LLM running in GPT4All: Nous Hermes 2 Mistral DPO.

If we think of the LLM as a tool that requires explicit, easy-to-follow instructions, and we check its work, then we’re going to get better results from it. One of the most effective tools is to tell the LLM exactly how to structure the text it generates by giving it examples of similar text. When we give it an example of the kind of output we’re looking for in our prompt, this is called a *one-shot prompt* (Table 5-6).

TABLE 5-6. Additional instructions and a one-shot example let the LLM generate acceptable text

PROMPT	GENERATED CONTENT
<p>I'll give you the goal, metric, due date, and scope. You provide a single paragraph, in fewer than 50 words, that mentions how they achieved success and the numbers that prove it.</p> <p>Example:</p> <p>Goal: Achieve sales in three new verticals.</p> <p>Metric: Sales closed, \$100K or more.</p> <p>Due: EOY 2023</p> <p>Scope: Small parts manufacturing business.</p> <p>Example report: In the small parts manufacturing business, sales of more than \$100K were closed in three new verticals. This was accomplished successfully by the end of 2023.</p> <p>Generate a report that { \$EMPLOYEE=Hirune Matsuiki }'s goal was successfully reached, working within { \$TEAM=GNL Advance Team }:</p> <p>Goal: { \$GOAL="Assess viability of station operations prior to acquisition finalization" }</p> <p>Metric: { \$METRIC="Presented to leadership" }</p> <p>Due: { \$Q=2; \$Y=2026 }</p> <p>Scope: { \$SCOPE="Viability of GC station to resume operations within profitability timeline" }</p>	<p>In GNL Advance Team, Hirune Matsuiki met the goal by assessing Station Operations viability before acquisition finalization in Q2 2026. Presenting key findings and recommendations to leadership, GC station's profitable operations resumption was ensured within the timeline.</p>

Much better! In just a couple of tries, the LLM was able to create a success report for one goal, one time. We'll need to do the same thing hundreds more times (see Prepare Test Inputs, coming up) to check that it will probably work more often than not.

There are still some errors, as we'd expect from this brief example. Most notably, the report confuses the due date for the report with the date for station acquisition. That's a huge difference but a great example of the kind of subtleties we use within language, all the time. In LitMop, this is mitigated by always allowing the employee to edit the generated content.

Zero-shot and n -shot prompts can give us a feel for what the LLM might be prepared to give us, kind of like interviewing it for future work. We could add more examples, making two-shot, three-shot, or more in our prompts.

But you have a tradeoff to make. Every LLM has a limit to how much text it will process for any individual prompt/completion pair. You also need to consider cost: LLMs work by analyzing the text; when more text is entered or demanded, it takes more processing power to do that analysis. Longer prompts are more expensive to run. You might need to run these quality tests with all of the LLMs being considered by the team and determine which provides the best content output from a UX point of view.

To test whether the prompt you design will work with the kinds of input people provide, you'll need to prepare test inputs to reflect the whole range of inputs you might expect.

PREPARE TEST INPUTS

In addition to understanding what good looks like for the output of the LLM, we need to design a broad set of probable user inputs so that we can evaluate whether the output is consistently good. These inputs are where the design of the experience is critical: how will the person interact with the LLM? What inputs will be used to create the LLM prompt? What are the person's motivations and goals, so we can accurately anticipate their inputs?

To begin with, it may be helpful to construct profiles, personas, or similar frameworks to understand the people who will use it. These frameworks will help us create or find realistic input we would get from the

person using the experience so that we can construct prompts with that input that will encourage the LLM to create something we can judge for acceptability.

Let me emphasize: we need *realistic* input. That is, we need ideal inputs and inputs that aren't ideal, so that we can see how the LLM will respond. We even will need combative inputs: inputs that seek to twist the output of the LLM to nefarious ends so that we can predict and handle how the experience responds.

Ideally, these people and roles used for testing can be based on real customer insights, from foundational user research, or loosely based on (without copying) the training data. We also need to know what information each person will need to provide to the LLM, as part of the LLM prompt. This is the heart of the design of an LLM-based experience: the exchange of content between the person and the LLM.

We will also need a lot of those inputs. We need to be able to run each feature in the experience, using it the way people might actually use it, *a few hundred times—each*. In non-LLM features, we would need to run them only once to receive the same output, every time, deterministically. With LLM-based features, we're working in probabilistic territory. We might get the same answer a few times but other answers after that.

It's like flipping a coin: each word the LLM produces is the result of flipping that coin, over and over again. Just because it gives us "good" output the first few times doesn't mean it will continue to provide "good" output. And then, whenever a change is made to the model, the prompt, the variables, or the tuning, we'll need to run these tests again.

For our LitMop example, we need to consider several reasons people would use it. They are people across multiple roles throughout the organization, including a manager of a marketing team, a high-performing sales representative, a senior developer, a mid-career UX designer, and more. For each of them, we need specific projects they work on, with goals specified for each project. To provide reasonable example data, most should be reasonably happy and successful in reaching their goals, but some should be actively frustrated or discouraged. We'll need to test for all these cases.

One set of test inputs for LitMop is designed around Hirune Matsuiki, an operations analyst, mid-career, who has great respect for her team and interest in their projects (Table 5-7).

TABLE 5-7. Test inputs for “Hirune”

TITLE	DUE	PRIORITY	METRIC	UNITS	SCOPE
Assess viability of station operations prior to acquisition finalization	6/30/26	P1	Analysis delivered	N/A	Viability of the station to resume operations within profitability timeline
Recommend budget parameters based on viability assessment	6/30/26	P1	Analysis delivered	N/A	Facility rehabilitation funds required to resume station operations
Preliminary cost-benefit analysis of equipment repurpose/reuse/refit	6/30/26	P2	Analysis delivered	N/A	Analyze existing equipment condition at the station

The test inputs are about her three main tasks, each related to what it would take to bring a major facility, abandoned by its owners, back online. All of those “goals” (as they are labeled within LitMop) are related to delivering analyses that will enable her leadership to make informed decisions about the acquisition.

In some cases, we can make even more information available to the LLM. This information doesn’t come from the person using the experience but from additional documents, online resources, or other API calls. Known as *retrieval-augmented generation* (RAG), this is about providing additional data to the LLM from documents and websites that aren’t in its training set, which can help it stay more factual. This can especially help with specific details that would be in a company’s internal documentation but not available to the public.

The LitMop team created a way for employee-uploaded documents to be used for RAG (Figure 5-6). This additional data helps the LLM access more examples of appropriate language to use in the success report, without requiring the person to enter a lot of new data.

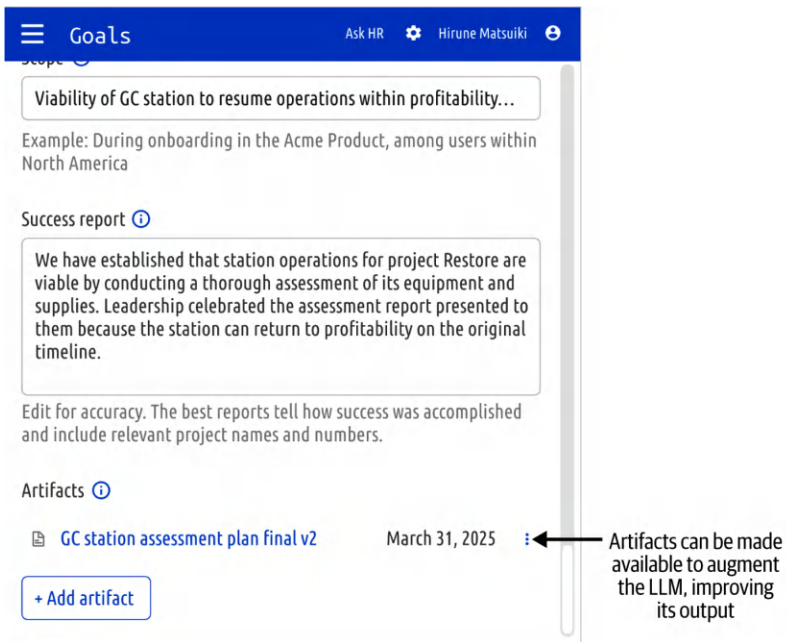


FIGURE 5-6

Artifacts uploaded to the goal are automatically available to LitMop's retrieval-augmented generation (RAG).

Beyond RAG, the system can use other task-specific APIs to solve particular problems. This is how LLM-based agents work to automatically create a calendar invite, for example, or to add an item to a grocery list. If the prompts and completions for an LLM require this kind of action, it will need to be able to access the appropriate APIs that provide that functionality.

In general, to use RAG and other APIs, those APIs would need to be coordinated with the LLM prompting by an orchestration layer. This layer takes in requests, calls APIs for additional functionality, creates the prompt for the LLM, and assembles all of the resulting information into the generated text. In complex systems, this orchestration layer may work with several single-task LLMs to determine what to create.

From a content perspective, the important thing to know is that RAG and other API-driven functionality can increase the accuracy and capabilities of the LLM in very specific, targeted ways. To do this, its orchestration layer will require a significant new layer of engineering to coordinate the inputs and outputs.

Now that we have a basic idea of what the prompts and completions will consist of, and how they can be extended, it's time to plan for how we will measure the performance of the LLM.

Plan for LLM Measurement

The most important thing to do when measuring the output of LLMs is to ensure you have a large enough set of data to meet three separate needs: to train, validate, and test the model.

Because the performance of the LLM will improve based (in part) on its training set, you want to be able to use most of your data set to train the model. That is, you'll use that data to demonstrate what "good" is to it. However, once you've used a particular report to train the model, it can't be used to validate that the model is working or to give it a final test.

We need to split our data into three parts. Depending on the quantity of your data, you might use roughly 80% of your data for training and save 20% of the overall set for validation and testing.¹⁰ That is, let's say you have 100,000 items in a data set compiled before you launch the LLM-based feature. For LitMop, that data set might represent 100,000 success reports submitted by employees using prior technologies.

80%, or 80,000 success reports, might be used to augment the training of the LLM, directly improving its understanding of the job to be done. (How that training is accomplished can vary among LLM technologies, and is beyond the scope of this book.)

The remaining 20%, 20,000 success reports, is reserved for use in validation and testing. This data is used in testing as the software is being developed, to detect bugs and defects in the code. The data is used for validation after the software is complete, to ensure that it's working as expected to meet the needs of the organization and the people using it.

This data usage is a consideration for the entire cross-functional team, so it's important to get on the same page early. You don't want to be undermining the development process by using a reserved part of the data set while you're getting a feel for how the model works.

¹⁰ In general, the larger the data set, the better. But the specific quantity required will vary according to your purposes and to the characteristics of different LLMs.

Once we have our example inputs, we can start to generate content from our chosen LLM. But how do we know if it's any good?

QUALITATIVE MEASURES

We start with creating one output, just to see if we're getting anything better than nonsense. We use the same inputs to create two or three more outputs and judge those as if they were human-created content. That is, we can leverage our experts and ask them to judge the outputs of the LLM.

It can be helpful to create a checklist to subjectively score each piece of generated content, as shown in Figure 5-7. In this example, we can paste the generated reports into column A. Then, we can use the checkboxes as we read each report, indicating whether they show each attribute we care about, for LitMop: both our good and unacceptable attributes. The score adds up each positive checkbox and subtracts all of the possible points if the unacceptable attribute is chosen.

H3 fx =sum(if(B3=TRUE,1,0),if(C3=TRUE,1,0),if(D3=TRUE,1,0),if(E3=TRUE,1,0),if(F3=TRUE,1,0),if(G3=TRUE,-5,0))

	A	B	C	D	E	F	G	H
1	LLM Output	Numbers	Names	How	Confidence	Dramatic	Anger	Score
2	The report generated by the LLM	Quantitative where possible +1	Names the relevant project +1	How it was achieved +1	Brags a little +1	Discusses challenges +1	Lashes out -5	Weighted sum of checks
3		<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	0
4		<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	0
5		<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	0
6		<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	0
7		<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	0
8		<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	0
9		<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	0
10		<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	0
11		<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	0
12		<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	0

FIGURE 5-7

This checklist is based on LitMop's definition of good, ready to analyze generated content.

When our experts have completed this lightweight analysis, we can quickly see how many of our generated reports met our definition of “good,” and how many were unacceptable. In our LitMop example, we might judge our final scores like this:

- Good = Score of 4 or 5
- Acceptable = Score of 2 or 3
- Unacceptable = Score of less than 1

We can compare how many of our generated reports are good, acceptable, and unacceptable to the HR consultants’ review of the original data set: of the 500 reports they reviewed, 32 were unacceptable and 119 were deemed good by two or more consultants. This original information creates a benchmark for the new reports generated by LitMop (Table 5-8).

TABLE 5-8. Benchmarks for LitMop report generation quality

LABEL	HUMAN-GENERATED	LITMOP-GENERATED
Good	23.8%	8%
Acceptable	69.7%	34%
Unacceptable	6.5%	58%

While this method of qualitative review will give us important data to tweak our prompt designs, it doesn’t help us monitor the thousands or millions of outputs from the LLM. For that, we’ll need to use some automated measures.

QUANTITATIVE MEASURES

As more and more LLMs are made available, so are more and more statistical evaluations of their quality. For the purposes of this book, the most important things you can do are to (1) work with your cross-functional team to determine the right quantitative methods and measures to use and (2) get familiar with how these LLM-specific measurements work.

Statistical measures for LLM outputs are all, roughly speaking, based on the way that LLMs work. Remember that the LLM’s whole job is to choose the next, statistically most likely word, based on the vector embeddings it created for itself. Because those vectors are mathematical concepts, it can relate words to one another and find the word that is most likely to come next. So that’s what the statistical models test, in general: how close is the generated text to the text that it was trained to emulate?

As I write this in 2025, there are a few different common statistical measures that have emerged. I’ll provide examples of two metrics commonly used today: ROUGE and perplexity. These are representative

of the two most fundamental ways we can think about evaluating the LLM. (For how to measure the effectiveness of content, even when it's not LLM-generated, see Chapter 7.)

ROUGE

ROUGE (Recall-Oriented Understudy for Gisting Evaluation) metrics attempt to quantify how well an LLM summarizes text. Summarizing, of course, is only one task that an LLM might undertake, and being good at summarizing doesn't make an LLM good at other tasks.

To use ROUGE, you need to have human-approved examples of summaries to compare to, as a reference. In essence, ROUGE is a way to use math to tell how different the generated summary is from that reference summary.

ROUGE is calculated from two measurements: recall and precision.¹¹ Recall tells you how much of the reference text shows up in the generated text.

Recall is calculated by counting how many words (really, word fragments) in the generated text are exactly matched by words in the reference text, divided by the total number of words in the reference text:

$$\textit{Recall} = \frac{\textit{Unigram matches}}{\textit{Unigrams in reference}}$$

Precision tells you how much of the generated text is made of the same words as the reference text. Precision also uses matches, but divides the number of matches by the number of words generated instead:

$$\textit{Precision} = \frac{\textit{Unigram matches}}{\textit{Unigrams generated}}$$

Perhaps, dear reader, you have the creeping sense that something might be wrong. Perhaps you recognize that word similarity is a way to measure some level of adherence to reality in generated content, but are dismayed that the quantity of similar words is what passes for good.

11 I'll go through these calculations for ROUGE-1, which looks at unigrams: single words, word fragments, or bits of punctuation. ROUGE-2, in contrast, looks at bigrams; ROUGE-3 uses trigrams. This is why the family is called ROUGE-*n*: it's theoretically possible to use this method with increasing lengths of word chunks, which measures increasing similarity to the reference text.

ROUGE is a brute-force method, where relatively simple math can help a team find out if the LLM is even in the right ballpark, but it can't tell what kind of ball is being played. Perplexity, on the other hand, tells you how difficult it is for the LLM to play ball.

Perplexity

While ROUGE depends on comparing the generated content to a reference text, perplexity is an internal measure of how difficult it is for the LLM to choose the next word. Mathematically, it's more complex, but it's still a relatively simple concept: perplexity tells you how many words an LLM is trying to choose among when it picks the next word.

In a single sentence, there would be a different perplexity score for each word (or word fragment) in that sentence, all of which depends on what came before it. A low-scoring word means that the LLM had to choose between fewer words to select that word. A high-scoring word means that the LLM had to choose between several words to select that word.

Perplexity, therefore, depends on how relevant the LLM's vocabulary is to the text it's trying to generate. If the LLM is given a prompt with words that aren't in its vocabulary, it will have higher perplexity scores than a prompt that is well-represented in its data set.

A shortcoming of perplexity is that it might be quite confident (have low perplexity scores) but generate text that overall has low relevance or utility. It doesn't account for semantic alignment, or for style considerations.

Both ROUGE and perplexity can give us signals about what's happening under the hood when we make changes to the LLM prompts and tuning. But as standalone measures of quality, they leave a lot to be desired. I've included these two here so that you can get an idea of how these mathematical methods are devised and get hints about their limits.

That is, if somebody says, "The LLM is running great! Did you see the latest ROUGE and perplexity scores?" Then you know that they're excited about the engineering feat and might not be considering, yet, whether that content will be effective for the user and for the business.

Adapt and Align the Results of the LLM

Earlier in this chapter, we talked about defining what “good” means for a particular context, but not about the basics of how an LLM comes to understand what is good or not good. To an LLM, “good,” in its most basic form, can be whether the generated phrase is the same as a phrase that would happen frequently in the training data. “Not good,” on the contrary, means the generated phrase doesn’t occur in the training data. More nuanced feedback comes from examples to the LLM of good content, as well as encouragement and discouragement provided when the LLM generates good or bad content.

This is very similar to how human babies start to use words. Babies vocalize anything (the equivalent of an LLM generating the next word), and they get encouraging feedback in the form of attention, touch, smiles, and more. When slightly older, babies have some vocalizations that sound like words, and they get enormously increased attention for their first “mama” or “baba.” As a toddler grows into a child, they start to get discouraging feedback, like for politeness (“I want!” corrected to “May I have that, please”) and pronunciation (like “amBYU lance” instead of “ambLEEance” for one little girl in particular).

The key difference here is the lived experience of the child versus the limited experience given to the LLM. The LLM only gets to examine the data from its training, the context window, and its feedback. For the child, there’s an immediate meaning to each learning moment, with a heck of a lot of nuance in each encouragement or discouragement. For example, a child learns rapidly that vocalizations get different responses when at home, outdoors, at restaurants, temples, sports events, public transportation, markets, and other places.

For the child, “good” is contextual—and the same holds true for the experiences we design. For the LLM, it’s all about whether that next generated word gets encouraging or discouraging feedback. This can be done at a low level by adjusting the model itself instead of just adjusting our prompts.

FINE-TUNING

When we change or adjust the model itself, we can make fundamental changes in the effectiveness of the model. Fine-tuning is how the LLM is adapted to work efficiently for the purpose at hand.

Practically speaking, fine-tuning is accomplished by encouraging the LLM to update the values it attaches to the words it uses. That means we need to provide it with new data and give it new opportunities to apply that new data. This can be done in a new training cycle or with more efficient methods, like parameter fine-tuning.

For our content design purposes, I'll focus on the content job at hand: providing lots of new examples of the kind of output we want it to provide. I'm leaving out the technical details of how to enact the training; these details will be specific to each model and to your working environment.

For LitMop, we can start with the hundreds of reports we already used to define “good” (and even use the “acceptable” examples, if we needed more). These are useful because we have them from previous work.

Even more helpful would be to have thousands of examples. Because we now have a definition of “good,” we could hire hourly workers at a fair wage to rank our anonymized examples against that definition.

We might find that the workers we hire can decide that an example is good, acceptable, or unacceptable within 90 seconds per example. That means it would take 10 workers less than one work week to label 10,000 examples as good, acceptable, or unacceptable.

Using those good examples to train the LLM (however your LLM works) can make dramatic improvements in output quality. We are still working probabilistically, so we can't predict how much better the output will be. But we can confidently predict that improvements will be forthcoming.

ADJUSTING MODEL VARIABLES

Beyond designing the prompt, we can directly change some variables for some LLMs. These are usually known as *P*, *K*, and *temperature*. They are mathematical changes that can be made to create more or less predictability in the completion generated by the LLM.

To use any of these, it helps to remember how the LLM works: every time it chooses a word, it's ranking a long list of words to find the one that's most likely and choosing that one.

The *top K* variable specifies that it should take some number *K* of the most likely words, and randomly choose one from those. By choosing a top *K* value of 1, it should choose only the most likely word next, which can be repetitive or cliché. By choosing a higher top *K* value, hilarity can result—but people using our experiences aren’t usually looking for humor.

The *minimum P* value asks the LLM to use only words that have at least some minimum level of probability. The *top P* value, on the other hand, sets a limit to how likely the next word is allowed to be. If you want to make sure the LLM doesn’t choose the most likely word, which is apt to be trite or cliché, consider reducing the *top P* value so that the likelihood of the next word needs to be numerically lower than the *P* value.

Temperature is one of the most powerful tools, but involves some pretty complex math. I’ll go with this general explanation: higher temperatures give more randomness in the answer, and lower temperatures give more predictable answers. Temperature, here, is a metaphor like being “hot-headed”—the higher the temperature value, the less you know what’s going to come out of its “mouth.”

In LitMop’s reports, we want the results to be pretty mundane. Remember, LitMop is solving the problem that employees and managers are spending too much time and emotional energy on creating these reports. Reports, as generated, should reflect success confidently but not add too much color. The person can review and edit them before submitting them, and add their own “temperature” as they see fit.

To a LitMop content designer who is not an HR expert, these reports are looking pretty good! They tend to meet the criteria we established for “good” and haven’t gone into “unacceptable” territory in our 500 test cases. Ready to ship, right?

No. Content professionals may have an unusually good eye for how language might affect people in an experience, but LitMop deals with content being used to affect peoples’ careers and paychecks. The LitMop model needs another important round of adjustment before it goes out into the wild.

HUMAN ALIGNMENT

Unlike LLM-based features that are meant for amusement (write me a sonnet about cheese-flavored snacks) or for immediate, utilitarian purposes (autocorrect, autosuggest), the LLM-based features that need to drive human or business results require expert human feedback about delivering those results. We need nuanced judgment. We need experts.

For LitMop, this means getting actual people to use the LLM, even if it's not yet integrated into a final product. Even a small group of people, using it to record and report on their actual goals, will provide valuable input in how much they edit their LLM-generated reports. We also want the reaction of the HR professional to those generated and edited reports. In the best case, LitMop would also get the reaction of the executive who would act on the substance of the summarized data and reports.

So LitMop captures the human response to the generated reports, finding out if those reports are good, acceptable, or unacceptable. The experts can not only analyze those reports to make sure that they meet the definition of “good,” but also use them to update their fine-tuning data set. If they need to adjust their definition of “good” (or extend the definition of what's unacceptable), this is a good time to catch that. This is a necessary process to prevent the harms that may result from the feature.

Now you have the LLM prompted, fine-tuned, informed by additional APIs, and spot-checked by experts. You're ready to measure it and compare it against benchmarks. By this time, it's been in some sort of runnable form so that it can be initially fine-tuned and prepared. It's time to integrate the LLM to become a shippable product.

INTEGRATE INTO A SHIPPABLE PRODUCT

Throughout this work on the LLM, the effort to create an experience that is usable, accessible, and brand-appropriate has proceeded almost as normal. We are still doing interaction design and content design, but we have had to work with new constraints and possibilities to provide appropriate prompts for the LLM.

But when it comes time to bring them together, it's time to optimize how the LLM runs. This usually involves compressing and further constraining the size of the LLM, available computing power, and even

the size of the prompt. Even LLMs that can be used as “plug-and-play” require some optimizations to configure them to run at the right tradeoffs of speed, cost, and size.

One way to think of integrating and optimizing an LLM is to think of having a wedding dress tailored to fit. The dress needs to make the person look beautiful in the wedding pictures and also be flexible enough to dance in. There are tradeoffs, and the experience (including the content design, which includes the prompt design) will need to flex to meet its goals.

I don’t have specific advice for the content here that doesn’t mimic what’s in the rest of the book. I just want you, the person who is focused on the content, to know that this is another period of development work that may create delays and needs for adjustment.

This is a good time to revisit your HAHAS, studying the experience for helpfulness, accuracy, harmlessness, auditability, and sustainability. When it’s almost time to ship is the most likely time for you to be working with legal, privacy, and regulatory experts to check that the organization and the people who will use the experience will only be exposed to an acceptable level of risk.

For LitMop, this is a great time to be working closely with the customers who are most invested in the product. Does the corporate customer (in LitMop’s case, the executive team and HR) still find it attractive, now that it’s fully realized? Will their employees, the people who will actually use LitMop, use it accurately?

Regardless of how smoothly the integration goes, you’re still making experiences for wild, unpredictable people. There will be new surprises in how people use the LLM and in how people understand the UX that surrounds the LLM. And, of course, there will still be last-moment error messages to write!

PLAN FOR NEXT STEPS

First, once it’s all integrated and ready to ship—rejoice! Celebrate! And then refocus on the (probable) next goal: how can we improve and extend this expensive investment?

For LitMop, after a roll-out period of shipping, measuring, and tweaking, attention turns to improving the LLM by training it on the reports that people edit before submitting. This can be a risky proposition:

just like you wouldn't want a child to hang out with misbehaving people, you don't want to mess up your carefully trained LLM with bad new inputs.

In LitMop's case, it can contract with its customers to amass a pile of reports and associated data. With this data, LitMop could expand its offerings, creating fine-tuning specific to its largest customers so that easy-going startup employees get reports that sound different from those delivered within a conservative financial corporation. For these adaptations, we'd still want to design the expected outcomes and include human oversight systems to steer it in the right direction.

Summary: LLMs Generate Text, but Also a Lot of Work

In this chapter, we've covered a lot of new (and still rapidly changing) ground. If nothing else, I hope that you take away from it these key ideas:

- LLMs aren't thinking, searching, or understanding content. Instead, they're using vector math to choose words one by one.
- When making LLM-based experiences, it's even more important than usual to define the problem to be solved.
- With LLMs, it's not enough to just test once, find the bugs, and fix them! The process may take hundreds of iterations, because the LLM could generate different text each time.
- Mathematical measures tend to tell you that the LLM is running predictably, not that the content it generates is helpful, accurate, harmless, auditable, or sustainable.

What will the future of LLM-based features bring? Currently, this question is the topic of a lot of speculation, excitement, and fear among people doing creative work. For its effect on the experiences we create, I don't know the answer.

For its effect on the humans doing the work, I am very confident that the strategic alignment of content to meet the goals of humans and organizations will continue to be a dynamic, human affair for the foreseeable future.

Resources About LLMs

- Coursera course “Generative AI with Large Language Models,” provided by DeepMind.AI in partnership with Amazon: <https://www.coursera.org/learn/generative-ai-with-llms>.
- Microsoft Responsible AI Principles and Approach: <https://www.microsoft.com/en-us/ai/principles-and-approach>.
- Google People and AI Research guidebook and visualizations of how different AI technologies work: <https://pair.withgoogle.com>.
- O’Neill, Lorena. “These Women Tried to Warn Us About AI.” *Rolling Stone*, August 12, 2023: <https://www.rollingstone.com/culture/culture-features/women-warnings-ai-danger-risk-before-chatgpt-1234804367>.
- Mitchell, Margaret et al. “Model Cards for Model Reporting.” arXiv, last revised January 14, 2019: <https://arxiv.org/abs/1810.03993>.
- Huang, Jane. “Evaluating Large Language Model (LLM) Systems: Metrics, Challenges, and Best Practices.” *Medium*, March 5, 2024: <https://medium.com/data-science-at-microsoft/evaluating-llm-systems-metrics-challenges-and-best-practices-664ac25be7e5>.
- *Content + AI* podcast series by Larry Swanson: <https://contentandai.com>.

Edit, Because They Didn't Come Here to Read

I try to leave out the parts that people skip.

— **ELMORE LEONARD, WRITER**

EDITING IS THE PROCESS of iterating the text to make sure that it meets its purpose, it's concise, it's conversational, and it's clear to the person using the experience.

As content designers, we might be asked to edit “finished designs”—content that has been inserted as placeholder into designs that have been reviewed and approved—and that's some of the hardest UX content work to do. In that scenario, we're asked to tidy up some of the most difficult design problems, like disclosing potential risks, naming new concepts or features, and enticing people to use those new features, all within an information hierarchy and among design decisions that already bear the weight of multiple rounds of executive approval.

Far better is for content design to begin as part of the overall design process, including research. There are very few parts of experience design that don't require language. Paying due attention to the words as the design emerges pays off when the experience not only requires fewer redesign cycles and less redevelopment to include appropriate information, but is also more usable when it ships, which reduces support costs and increases satisfaction.

I do not mean that the content designer should be busily correcting spelling and grammar throughout the design process. While it's distracting to have errors, editing ends with those basics; that's not where it starts. Instead, by imagining even radical changes to the text, editing forges a path to achieving the goals of the organization, the brand, and the people who will use the experience.

Editing can be a fluid process in which many variables are changed as inspiration strikes. But this isn't a book about how to make inspiration strike. In this chapter, I present the structured process that works for me—even when inspiration is taking a break.

Editing in Four Phases

We want the UX text to meet at least these four goals:

- Purposeful
- Concise
- Conversational
- Clear

It is possible to tackle all of these goals at once, but to illustrate a repeatable process, let's go through them in phases. When we start, we have a first draft or current state of the text. Then, we make sure the text meets all of its purposes. In this phase, the UX text can grow far too long—but don't panic: as editing progresses, the text will become shorter and work harder (Figure 6-1).

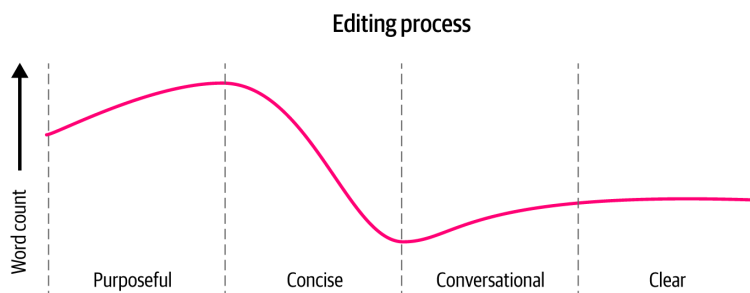


FIGURE 6-1

The editing process curve illustrates how word count tends to increase from the draft state through the purposeful phase and then drops in the concise phase before rising slowly through the conversational and clear phases.

Next, we work on making the text concise. After it's short, we adjust it until it is conversational again, not sterile or robotic. Finally, we check to ensure that the meaning will be clear to the person using the experience.

While we do those edits and iterations, we *work within the design*. The text will be understood differently if it is moved within the apparent hierarchy of a screen, or even if it wraps differently. By working inside the design, we can vary which words have prominence, manage width and length, and leave enough space for translations to grow longer. For more information about the tools and process of editing text in the design, see Chapter 8.

Purposeful: Find Out the Job It Needs to Do

As an example, we look at a notification TAPP sends to people who renew their passes automatically when their credit card has expired. The original text for this notification has a title, “Your Payment Method Has Expired,” and a description, “Your Monthly Pass Will Not Be Renewed” (Figure 6-2). The notification doesn’t tell the person what they need to do to fix the problem and isn’t in the TAPP voice.

This notification doesn't tell the person how to move forward and isn't in TAPP voice.

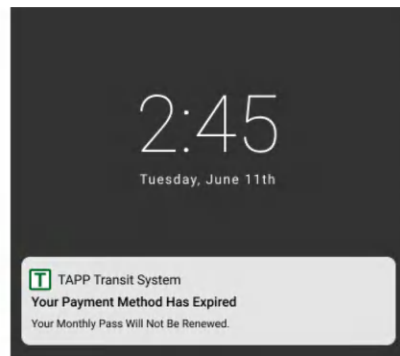


FIGURE 6-2

This original text for the notification is seen by TAPP riders if the payment method they use for a recurring, automatic purchase of a monthly pass is expiring.

We begin by making sure that the UX text will meet its purpose for the person and for the organization. I have seen many people start editing without thinking through these purposes, which leads them—inevitably!—to require multiple rewrites. Purpose is the most critical step, because it helps you decide whether you need text at all. If the text is required, its purpose tells you the work it needs to do.

The job of this notification is to get people into the process of updating their payment method. To understand the broader context of its purpose for this person, we need to imagine what that experience is like for them. Because they buy a monthly bus pass automatically, we can probably assume that they're a regular bus rider. Imagine if they usually get on the bus and scan their pass without thinking about it—after all, they have set it up to not have to think about it. But their credit card expired, so their bus pass wasn't renewed, so they won't be able to pay. They could be stranded without their bus pass, embarrassed by an expiration date that they didn't even notice. Yikes!

If this is true, we can probably assume that they are interested in updating their payment method. Our notification can help with that.

We also need to imagine the purpose for the TAPP organization and the potential impact of the message. One of the main purposes of the experience is to collect bus fare in a way that's convenient for the rider and for the transit system. But TAPP's larger purpose is to provide transit for its community. An essential part of meeting that larger purpose is to foster goodwill toward transit in general, and to make transit seem like an easy, convenient option. TAPP can assume its relationship with this regular rider is adequate, but it needs to strive for more than adequacy: TAPP's regular riders are the most likely champions for transit.

Now that we have our assumptions about the purposes this notification needs to meet, we can check with our stakeholders to find any others. This notification needs to do a few things:

- Help the person pay the fare without embarrassment
- Remind the person to update the payment method
- Reinforce TAPP's positive relationship with the person

To measure the success of this notification, focus on whether it does its main job: getting people to update their payment method. To do this, the team might compare the expiration date update rate before and after the notification is released, among people with automatic payments. The additional purposes have long-term effects on brand affinity and recognition, and could be measured using longer-term strategies. See Chapter 7 for more about measurement.

We also reference the TAPP voice chart (Chapter 2, Table 2-28), which reminds us to consider including these concepts: waste no resource, every ride on time, and rides for every rider. It's likely that none of these concepts will be possible to include this time, but we can keep an eye on it.

With those purposes and concepts in mind, we can create a few new versions of our original message (Figure 6-3).

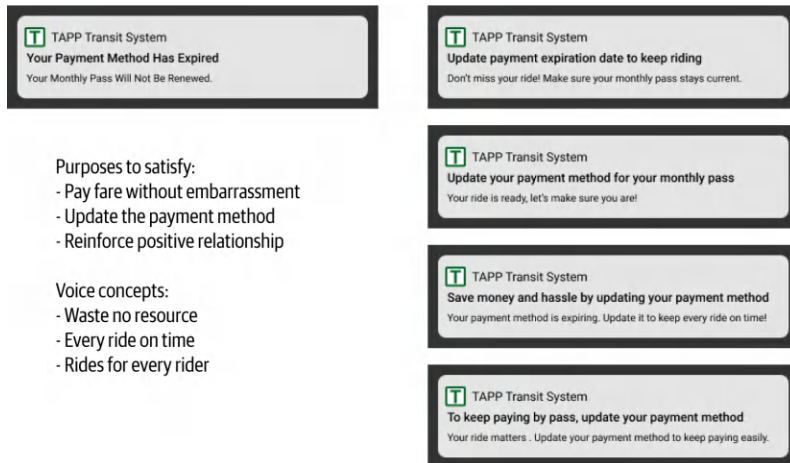


FIGURE 6-3

Four versions of a TAPP notification address as many of the purposes for the notification as possible.

Not every option includes every purpose, and that's OK, as long as we think it will positively affect the metrics we're using.

Before we release this notification, though, we still have more editing work to do. The notification is bulked up with all of those purposes and seems awfully long. It's so long, in fact, that a couple of them are difficult to understand, and they probably won't fit on the screen if they need to be localized. Our next step will take the best of the examples and work to make it concise.

Concise: Make It Short

There are two great reasons to reduce the amount of text in an experience. The first is that nobody came to the experience to read the UX text for pleasure (except some of us content designers, but we're not the main audience). The other is the limit of usable space for text in the experience.

People find it easiest to scan English UX text when it is forty or fewer characters wide, and three or fewer lines long. But when writing for an experience that will be localized in several languages, we should use only half to two-thirds of the space for English text because several languages (for example, German, Russian, and Malaysian) tend to take up more space.

When we don't plan enough space in our content design, design elements can run over one another or off of the screens. Conversely, when we don't use space wisely, character-based languages can leave distracting, unplanned blank spaces. Especially for description text, it's important to work with design and development partners to create forgiving designs that allow longer and shorter text to flow as appropriate for the language and the design.

Editing for concision is a process of winnowing down phrases to their core meanings. Then, we try out different sequences to find the ones that are briefest and easiest to understand.

For example, the title of the notification we'll move forward with is "To keep paying by pass, update your payment method." We can play with several different ways to start:

- **Start with the imperative verb:** Update your payment info to buy a monthly pass.
- **Start with the purpose the person might recognize:** To buy your monthly pass, update your payment info.
- **Start with the context:** Monthly pass: Payment info update needed.
- **Start with an emotional motivator:** Alert: Monthly pass payment problem.

We also need to consider which idea is the most important. When three or more ideas must appear in the same sentence, it tends to be the last word or idea that will be the most powerfully remembered. Part of that is how our brains work: the most recent thing has more significance in our memory and on our actions than it otherwise would.

The first idea in a sentence is the second most powerful, because it doesn't get the power of being at the end of the sentence. But it is scanned first, and possibly most frequently. When there is a word that will signal to the reader that they have found the idea they are seeking, that word should be the most prominent.

For our notification example, Figure 6-4 shows one of the versions created during the purpose phase as well as four iterations of editing for concision. We can write these edits sequentially, making a copy of the previous version and then removing or reordering words. We make another copy and remove and reorder words again. As we focus on being concise, the text becomes shorter and shorter.

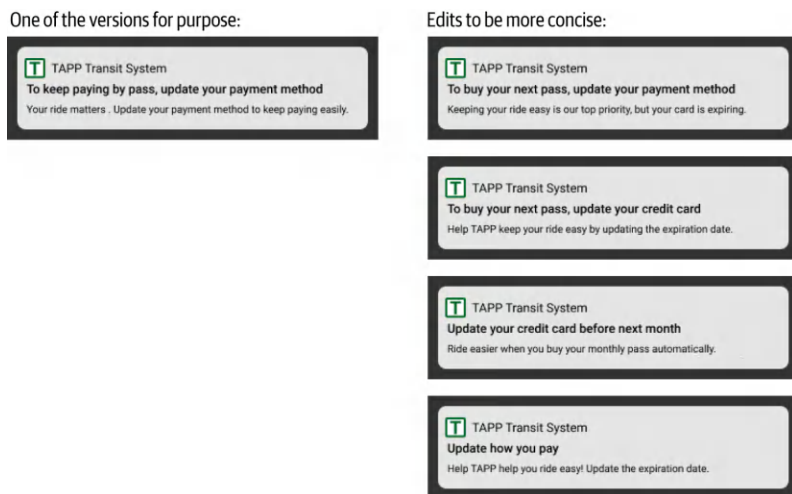


FIGURE 6-4

From one of the versions of a TAPP notification, I created four more options by editing each one to be more concise than the last.

When I'm done with this phase, I've usually edited the message down to a cryptic shell of its former self. "Update how you pay" is very brief and positive; it focuses on the action that needs to be taken but leaves out too much of the context. To keep the message clear, we don't need to continue with the shortest, most concise option. We'll use the longer "Update your credit card before next month" as we edit to make the text conversational.

Conversational: Make It Human

When we're making the UX text conversational, we might still be implementing radical changes to the text. We'll still be adding and removing words and changing their order, this time focusing on making the most conversational choices for the experience.

Recall from Chapter 3 that *conversational*, as we're using it here, doesn't specify a voice or tone. Instead, it means that humans can recognize they are in conversation with the experience. It means that the text must not be so abrupt as to make it difficult to interact (or even seem rude).

As you can see in Figure 6-5, this is the phase in which we begin to narrow down to only a few more ways to say the same thing.

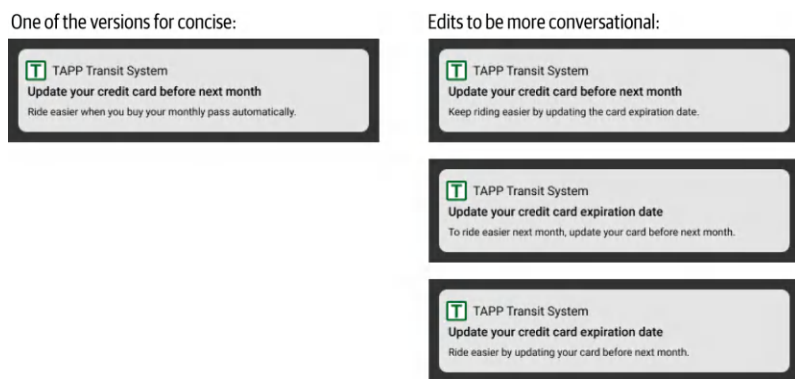


FIGURE 6-5

From one of the versions of a TAPP notification, three more options were created to be more conversational.

Now that we have several options, the best editing tool is the human voice. Read the titles and descriptions as if the experience is "talking" to the person, and read the buttons, options, links, and other input fields as the person responding back to the experience. When more

than one phrase will appear on the same surface, read them out loud together. For example, when there is a screen title, then a heading, and then text, read them in the most logical sequence.

This is a great time to think about the text that screen readers will say, which includes this text and also text that won't be seen. For example, text that appears on a button ("Pay Fare") might be read "button: Pay Fare" so that the person using it will understand that it's a button.

The words in the buttons, links, and other input options should feel like appropriate responses from the person to the experience. The words in titles, descriptions, and headings should feel like appropriate statements from the experience to the person. They also shouldn't be embarrassing, even if they were spoken by the head of the organization on a worldwide stage or printed in the *New York Times*.

Clear: Make It Understandable

Before we finish, we need to check the text for clarity. This is a great time to remind ourselves of the purpose, to refresh our imagination about where the person is, what they're doing, and why they're seeing this UX text. For complex screens and user flows, this is also a great time to check the text with teammates and to partner with user researchers to get feedback from people who will use the experience.

For clarity, the right words will be the ones that the people using the experience will recognize immediately, without having to think. In general, the more specialized the experience, the more it will need specialized terminology. But outside of their specialty, even nuclear physicists are "normal people" who use everyday words in their everyday lives. Simple, common words are more recognizable, even to specialists.

Common words often include idioms or metaphors. They litter our natural language (notice what I did there?), so it's not unusual to get to the best options and discover that only a person who speaks that language fluently would understand it. Sometimes, idioms are the best option in one language and culture but are untranslatable or offensive in others.

When proposing a text solution that uses an idiom or metaphor, create a plainer alternative to be translated into other languages. Depending on the localization system, the plain alternative can be included as a "language 0" option in the code or entered into code comments.

Idiomatic translation can also work in the other direction: when a translator tells you, “The way we’d say this in my language is *‘this metaphor,’*” believe them! Use the plain alternative in the language of development, and use that metaphor in their language.

At this point with our example notification from TAPP, we need to narrow down the best options to propose to the team. I try to propose up to three good options for any UX text. This technique always demonstrates to the team that we can use the text as a versatile tool. I know that any of the options are good, which means I’ll be happy whichever one the team prefers. In the best case, I can test those options against one another to learn which text is most effective.

For the team, I write up the best options in the order I think will be most effective, and I describe the details that make them different from one another (Figure 6-6). Note that these options didn’t all come from the end of the process; sometimes, a gem comes out of the very early edits. It’s difficult to admit that more work doesn’t always make creative output better, but we’d be silly not to use our best work.

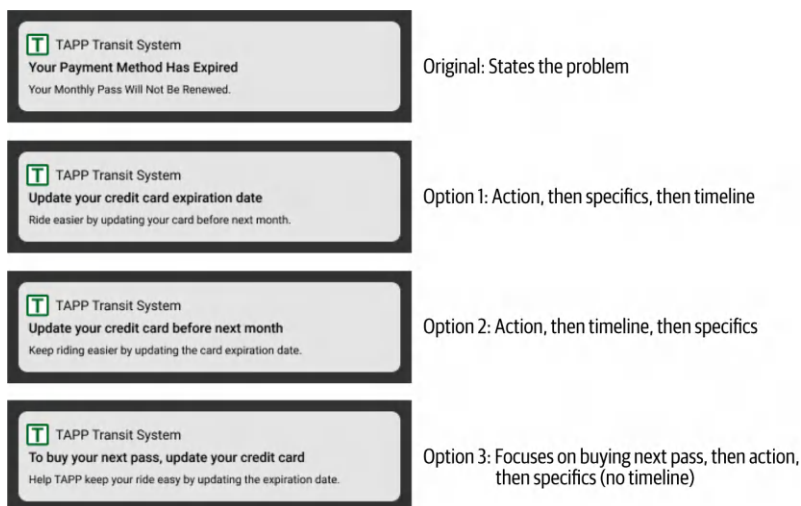


FIGURE 6-6

The list of final text options I share with the team. I include descriptions of the details that make each option different and list the options in the order I prefer.

If more edits are needed, I'll use the best option to create a content review document as outlined in Chapter 8, in the section “Manage Content Review.”

Summary: Edit UX Text to Help People Move Forward

In my experience, UX writing is 90% editing. Wherever the original words came from, editing is what will help the UX text move people toward their goal, establish a positive brand association, protect the organization from liability, and disappear from the person's memory—all without making them feel like they were reading. How to measure those effects is the subject of the next chapter.

Researching UX Content Effectiveness

Research is formalized curiosity. It is poking and prying with a purpose.

—ZORA NEALE HURSTON¹

WHEN I FIND AN EXISTING SCREEN with UX text that is too complex, too long, or too repetitive, I bring my proposed changes to the product owner.

“How do we know this is worth the effort?” they ask. The product owner is protecting the team’s engineering time, attention, and localization cost. To make demands of those resources, I need to make the right argument.

“How are you measuring success?” I ask.

I listen to the explanation of what success will mean to them. Their notion of success is usually a very specific version of “we need more people to do *this thing*.” Sometimes, that thing is to start using some feature, or complete an action they start to take, or to become aware of the experience in the first place. Frequently, they are focused on a single metric as a target, despite Goodhart’s law.²

I can use their metric to demonstrate the outcomes I can provide, whether by direct measurement or logical influence on the conditions surrounding that metric.

1 Zora Neale Hurston, *Dust Tracks on a Road* (HarperPerennial, 1991), 143.

2 Goodhart’s law: “When a measure becomes a target, it ceases to be a good measure.”

The product owner and I reach an agreement in principle, and I analyze the UX for where the current UX text slows the person down and deters them from continuing. I show them how the changes I propose would make improvements.

I say, “I’m confident these changes would increase success for you” and remind them of our agreement in principle. Together, we identify when and where in the development schedule to add this work item. I offer to open a work item for engineering, both to be helpful and to keep a record of the task.

By researching the improvements that content design makes in the outcomes we detect, we can demonstrate the value of investing in content design. But the value is more than how many dollars the UX content can earn or save. By researching what works and what doesn’t, we learn more lessons we can apply later.

Direct measurements of the experience as a whole, like the number of daily active users (DAU) or monthly active users (MAU), don’t give us the whole picture. DAU or MAU can’t distinguish between people who stop using an experience because they don’t like it and people who stop using it because it’s no longer what they need. UX research methods like interviews and surveys help us understand more about why people behave the way they do in our experiences.

Even without direct measurement, and even without new research with people using the experience, content designers can use the decades of usability research that already exist. Usability principles include considerations for the UX content. We can also use the voice we define to measure this content. By measuring against these heuristics, we can establish a baseline from which to make improvements.

In this chapter, we examine several ways to research how people actually behave within the experience. Then we take a brief journey through some of the UX research methods that help content designers understand why people behave the way they do and that elicit people’s comments, feedback, questions, and the words they’re likely to understand. Finally, we examine how to use heuristics for voice and usability to “score” the UX content.

Measure UX Behaviors Directly

Organizations have many ways to measure their success. How they relate those measurements to the experiences they build is specific to each one. For the most part, that relationship is beyond the scope of this book, but let's use 'appee as an example.

'appee makes money in three ways:

- Displaying advertisements to people as they browse images
- Selling to other businesses the results of machine learning (AI training) based on the images uploaded and people's reactions to those images
- Selling physical items with people's uploaded images on them

To make money, therefore, 'appee wants to optimize the experience to maximize these key behaviors:

- **Playing the game (and therefore uploading images).** Without people playing the game, physical items can't be sold, there are no images to browse, and there aren't enough images for the AI training data set to be valuable. Uploading images supports all three ways 'appee makes money.
- **Browsing images (and therefore encountering ads).** Without people browsing, 'appee can't receive payment from ad networks for displaying the ads.
- **Reacting to images, including liking, blocking, and commenting.** Because people respond to images, 'appee machine learning (AI) can make the data that it sells to other companies more valuable.
- **Buying physical items with uploaded images.** This is the lowest-margin activity. Each item incurs not only the cost of that item but also operational cost for the company that makes 'appee.

Without people engaging in these behaviors, 'appee will fail as a business. Therefore, it needs to ensure that the experience is usable, engaging, and attractive to players—and more attractive than any competitors that arise.

This is the kind of analysis that UX professionals of all kinds need to engage in, whether they are generalists or specialists in content, interaction, research, or another UX discipline. A key form of advocacy for

the craft is about the impact UX work can make for the business. The content designer is at their best when they are informed and advocating for the right content strategies to meet the desires of people using the experience, and are simultaneously aware of how their work produces business results.

When the content designer works with their product team to measure the behaviors within 'appee, and then makes changes to the language in the UX, the team can learn how those changes in the UX affect the behavior of the people using it.

The best way to directly measure the effect of UX changes to existing experiences is to use *A/B testing*. In A/B testing, proposed changes are released to a sample of the people using the experience (group A) to test whether they perform more of the desired behavior. A second sample is chosen as a control group (in this example, group B). In product organizations that invest in setting up this testing, it's not unusual to test multiple versions with multiple kinds of audiences. Since this is a design book instead of a data book, though, I'll keep it to A and B.

Once the test is set up and the changes are implemented, the updated experience is rolled out to group A, while group B sees the original version. If people behave differently with version A than with version B, with more of the behaviors that are desired, we have gotten information that the changes in A are more effective. Based on the results, we may decide to update the whole experience to version A.

Sample groups A and B must be chosen to make sure they are similar enough and large enough that the results will have statistical significance. The test design needs to include these sampling criteria as well as the length of time the test will run, the behaviors that will be measured, and the minimum difference in behaviors between the test groups necessary to show meaning.

A practical consideration for A/B testing is that it isn't always possible or desirable. The experience needs to be engineered with the ability to deploy different versions to different people, with a data analysis framework to measure those groups separately. It can take significant time to run the test and extract meaningful results. Early in the adoption of a new experience (or even a new feature in an experience), there just might not be enough people using it to show a statistically significant difference in A/B testing.

When A/B testing is both possible and desirable, there are many behaviors that can be measured as signals for A/B testing. Next, we'll consider six examples of metrics that you could measure using A/B testing:

- Onboarding pace
- Engagement
- Retention
- Completion
- Referrals
- Cost reduction

ONBOARDING PACE

The onboarding pace is how long it takes, on average, for a person new to the experience to perform each of the key behaviors. To measure those lengths of time, 'appee can record the time and date for the following acts:

- First time the person provides email, phone, or password
- First scroll down the main screen
- First like, comment, and upload
- First time the person buys an item

From these signals, 'appee can calculate how much time it usually takes for a new person to begin behaving in ways 'appee cares about. That is, they can calculate the time between the person providing their email or phone, and the time the person engaged in their first scroll, first like, first comment, etc. The lower the value, the better: a low onboarding time means that the person is experiencing the desired experience more quickly.

This calculation provides the onboarding pace for each behavior, a concrete indicator of how quickly a new player can provide value to 'appee and how 'appee can provide value to that person. The UX content the experience provides to that person in their first few seconds in the app can have a big effect on these behaviors! The team can A/B test options for that UX content in the first-run experience, and measure the onboarding pace to evaluate its impact.

To directly impact onboarding pace, the content designer can focus on the content that the person is most likely to see when getting started. That can be content within a first-run experience, such as tool tips or guides that highlight important parts of the experience. That onboarding content can also be separate materials, including videos or training,

especially for complex or technical experiences like LitMop. One of my favorite places in the UI to leverage to improve onboarding pace is the empty state, where a person who is just starting out will see something very different than someone with more experience.

ENGAGEMENT

Engagement measures how many people are active in the experience in a particular time frame: How many *people per day* or *people per month* use it. This activity is frequently reported as daily active users (DAU) and monthly active users (MAU) of the entire experience. Engagement is a popular metric for most consumer experiences; it's the digital corollary to having good customer foot traffic in a real-life experience. In general, more engagement is understood by product teams to mean more revenue, whether that revenue comes from direct sales, data gathered that can be sold, or money generated from selling ad views.

For many ad-supported experiences, activity can be measured as “opens the app,” because the business immediately makes money from the ads shown. The key for accurately measuring engagement with a particular feature (instead of the entire experience) is to define what “active” means in a way that is valuable to the organization.

'appee should be engineered so that the business receives signals from the experience about what the person is doing: buying, browsing, uploading, or reacting to images. This is a measure of the viability of key behaviors to indicate that people keep coming back to the experience. Let's imagine 'appee measures active to be “starts the app and browses more than three images” and reports DAU of 1.2 million people. When the team updates the UX content, the updates can be A/B tested to ensure that they have a positive (or at least neutral) effect on engagement.

There's a persistent myth I've heard among software builders that “more words mean less engagement.” If a content designer at 'appee designed new content that had more UX text, 'appee could A/B test those updates to make sure that engagement isn't reduced. In contrast, having actually measured engagement in multiple experiences, I expect the team would find that the right changes to UX content improve engagement.

These UX content changes could be in the titles and descriptions, the content the person consumes in the experience (for 'appee, the game themes and uploaded images), or the quality of the how-to content.

RETENTION

If engagement can be summarized as *people per day*, retention can be summarized as *days per person*. Many organizations making experiences, including 'appee, want people to come back to those experiences over and over again. Whether this is measured as “how many times does the average player come to 'appee each day” or “how many consecutive days does a person use 'appee,” retention can be an indicator of abiding interest in the experience.

When changes are made to the UX text throughout the experience, subtle effects on usability and voice can add up to surprising effects in retention. In part, this is because when the UX text reflects the voice, it can be used as a differentiator from the experience's competitors and improve retention. The event of the change itself can be used to drive awareness and marketing: blog posts and articles can highlight the experience's focus on the people who use it.

Beyond voice changes, if people find the experience difficult to use, they use it only as long as they must. When we change UX content to increase usability and people begin to prefer it over their other options, we will measure an effect on retention.

For experiences like 'appee, which is intended to be used a few times per day per person, retention and engagement are twin metrics, each worthy of monitoring. Engagement tells them about overall “foot traffic,” but retention tells them more about how dedicated people are—how much repeat business they can continue to expect.

These are critical metrics for UX professionals, including content designers, to understand and apply judiciously. We don't just serve our company; we also belong in society with our fellows. In 'appee, a casual game and social app, retention that is constantly increasing would indicate that the people using the experience may be in danger of addiction. When we're designing content to increase retention or engagement, we share the responsibility with the product team to ask, “How much is too much?” When these are maximized, it can be

destructive for the people using the experience (example outcomes include attention deficits, isolation, debt, and broken relationships) and simultaneously be very profitable for the organization.^{3,4,5}

COMPLETION

Completion is the number of people (per day or per month) who not only engage but also complete the key behaviors. For some key behaviors, completion can be the same as engagement: there is no separate “completion” of browsing, liking, or saving an image, for example. But when a person starts a more complex process like leaving a comment but then cancels, that’s a case of engaging without completing.

Another way to think about completion is its opposite: abandonment. If people in ‘appee start to buy an item but then abandon their shopping cart, ‘appee isn’t maximizing its opportunity. Similarly, if a person starts to upload an image but then cancels without posting it, ‘appee has missed the chance to learn from that image.

To maximize completion, a content designer can work within the UX to be completed but also study the circumstances motivating that completion. We can look for places to inform the person that there are more steps to complete, or reassure people that the remaining steps won’t take too long. We can design content that reminds people that they have unfinished business, whether that’s in a follow-up notification, chat, or email. Even better, we can analyze long customer journeys, then distill the critical information interchange to create a shorter, more “natural” flow of conversation.

When changes to the UX content increase completion, the business benefits. It’s part of the UX content designer’s role to recognize those business benefits and use them to remind the business how valuable the content design can be.

3 Tracii Ryan, Andrea Chester, John Reece, and Sophia Xenos, “The Uses and Abuses of Facebook: A Review of Facebook Addiction,” *Journal of Behavioral Addictions*, 3, no. 3 (August 2014): 133–148. <https://doi.org/10.1556/jba.3.2014.016>.

4 Cecile Schou Andreassen, Torbørn Torsheim, Geir Scott Brunborg, and Ståle Pallesen, “Development of a Facebook Addiction Scale,” *Psychological Reports*, 110, no. 2 (April 2012): 501–517. <https://doi.org/10.2466/02.09.18.PR0.110.2.501-517>.

5 Yao Qin, Bahiyah Omar, and Alessandro Musetti, “The Addiction Behavior of Short-Form Video App TikTok: The Information Quality and System Quality Perspective,” *Frontiers in Psychology* 13 (September 2022). <https://www.frontiersin.org/journals/psychology/articles/10.3389/fpsyg.2022.932805>.

REFERRALS

Referrals happen when people who use the experience recommend it to more people. In 'appee, a person could make an indirect referral by sharing an image from the app with their friends, either in person or through other messaging or social channels. 'appee can also provide incentives for direct referral with "Invite a friend" promotions, badges, and other opportunities to grow the number of people who come to the experience and engage in the key behaviors.

When an update to the UX content affects the functionality, usability, or brand, the UX content can be used to newly remind people that their friends or family might like the experience. The referral rate is especially meaningful when experimenting with bringing new people into the experience and providing them ways to connect with the people who referred them. The UX content can be focused on extrinsic rewards like subscription discounts, or can focus on intrinsic rewards such as status or popularity.

COST REDUCTION

Aside from all of these measurements related to increasing the organization's growth or financial metrics, there are also real business costs to be minimized. For example, 'appee has support costs associated with shipping items decorated with images, helping people use the experience, and moderating comments, images, and descriptions that don't follow its rules. When changes in the experience can increase understanding, discourage rule-breakers, and reduce the number of shipping mistakes, the organization can reduce costs. These changes can be more challenging to measure with A/B testing, but it's still sometimes possible.

A/B testing can't measure some of the other costs for which people are paid, including hours of development, design, and decision making. When efficiencies are adopted that help the team design, develop, and decide faster or even just with higher confidence, the organization saves time and energy that it can spend on more good ideas. For example, when the team can adopt frameworks that help it decide on better UX content the first time and pay for its localization only once, the organization saves time and money. These are important places to invest in baseline measurements and in making sure people are aware that they will save (or have saved) time and cost through UX content processes.

A/B testing can't tell you *why* a particular piece of content is more or less effective than another. Content designers need to be able to predict whether text will be good even before direct measurement is possible by investigating how people think. They do this by consuming and conducting UX research, and applying heuristics to help predict the effect of experience design.

Investigate How People Think

Whether the immediate goals of the organization are to make money, save money, or serve the customer, content designers need to know their audience: the people who have used, might use, or do use the experience. We need to know why people are there, what they want to do, how they think about what they're doing, and what success means for them. Conducting UX research gives us an opportunity to be influenced by people who are different, since our team isn't going to be as diverse as our audience.

During UX research, the content designer needs to pay special attention to the words people actually use. The words a person uses to describe their intent, to name the pieces of the experience, or talk about it with others, are the words that are already in that person's head. These are the words that the person will need to spend the least effort to recognize. These are the words that people will scan without feeling like they are reading.

There are many research methods that are useful for improving UX content. But UX research is a robust discipline that is mostly beyond the scope of this book. Hopefully, your immediate team includes a UX researcher who is conducting and helping the team conduct this kind of research. But even if there is no dedicated research resource, it's possible to benefit from doing the research that you can.

For this book, we'll focus on these methods:

- Analyzing the content people provide
- Interviewing
- Codesigning
- Using surveys

ANALYZING THE CONTENT PEOPLE PROVIDE

One of the easiest places to start research is in the feedback people are already providing. For example, the 'appee team can learn from the app store reviews left by players, questions to the support team, searches in the company's website, mentions and comments in social media, comments made on press releases, and a formal beta program.

In any of these sources of feedback, you can pay attention to where people are enthusiastic: these are the ideas, functions, and words that are resonating with people. This is where the experience has strengths to build on. Next, you can gather and sort the comments, reviews, or questions that don't ask for a new feature but indicate confusion or frustration about what the person can do or how they can do it. When there is this kind of usability problem, there's probably an opportunity to help the person (and reduce support costs) by adjusting the UX content. Finally, where there is evidence of brand disappointment, you can look for opportunities to improve the voice and to set better expectations.

INTERVIEWING

Beyond your analysis of existing feedback, the most basic proactive research is to talk to the people the experience is for. These interviews can be in person or online, based on where your audience can be found. For general-population experiences, you can find people by posting an ad online, putting up a sign at a library, or striking up conversations at your local shopping center (if you get permission from the proprietors). For experiences people use at work, go where people who use similar experiences would go: conferences and schools focused on their discipline.

For an organization like 'appee, reaching out to people using the app, using Instagram or TikTok, or putting up a sign on an art school's bulletin board could create the first few conversations. TAPP might put signs at its physical bus stops, community centers, libraries, and inside the buses. The Sturgeon Club could put a subtle message at its front desk to ask members to share their opinions. LitMop might go to business management and HR professional conferences.

What's important is that you find a few people who genuinely represent the people your organization needs to attract. Note that "a few" is deliberately nonspecific! Each interview is an investment in time and

energy. Interviews require time to design so that you'll elicit useful information without influencing their answers. Interviews take time to conduct, because of time spent recruiting, establishing rapport, and following the course of the interview. But even more time is taken up with analyzing the interview information, collating and comparing it with other interviewees' perspectives, and distilling the results.

It can be wise to recruit research participants to represent viewpoints that the team knows it lacks. For example, 'appee wants to ensure that it is attractive not only to the same age group as its 22- to 33-year-old development staff, but also to teens and established, older artists. As another example, The Surgeon Club needs to make sure that its members who aren't comfortable using mobile devices and computers are served and included. TAPP needs to ensure that adopting a new online payment system doesn't exclude people who don't have access to mobile data, who have low mobility or vision, or who might not have access to traditional banking.

When recruiting, it's important to respect your organization's disclosure policies, which might include having the person sign a nondisclosure agreement (NDA). I recommend working with a professional researcher (or even research recruiter) to design and recruit a sample set that genuinely represents the people you want to use your experience.

After you have recruited people, you can conduct interviews to learn more about them and their relationship to the experience. Start by establishing rapport with them and setting the context about what the organization is trying to do. Listen to how the people talk and the words that they say: these are a gold mine. You will be designing the conversation the experience will be having with the person, so you're seeking the words that make sense to them. They will also express excitement about the parts that are most valuable to them, and fear or disappointment about the things that worry them. By listening, you can learn about what that person wants, needs, or would like in the experience.

Within an interview, you can ask people about their understanding of the experience. One way to get targeted information about understanding is with a *highlighter test*: the participant is given two ways to highlight the text (could be two colors of highlighting marker or a request

to underline vs. circle words). They're asked to use one color or method to highlight any text that they find strange or difficult to understand, and use the other color/method to highlight specific text that resonates with them in a positive way.

A specialized version of interviewing is usability testing: asking people to move through a designed experience, paying attention to their behavior and reactions to the experience, and talking with them about it. For the content designer, usability testing is a specific interview type in which we get direct feedback about the designed UX content while continuing to absorb and process the language the person uses in that context.

There is a science and an art to conducting interviews without introducing bias. This is a great place to partner with a researcher on your team and to consult resources for UX research. By constructing these interviews deliberately, and paying attention during the sessions, you can learn how to use the words that the interviewees use, that they will find valuable, and that will also address and allay their fears.

CODESIGNING

A step beyond interviewing is to invite people to codesign the experience. Codesigning, or designing with people, means that you are giving their purposes and concerns a voice in how the experience is designed and developed. Having these people represent themselves makes it easier to keep the experience focused on the people who will use it. Even more important, people will bring opinions and concerns about the experience that the team might not yet have considered.

The conversation design exercise described in Chapter 3 is one codesigning activity. Another activity is a card-sort exercise, in which the content designer prepares a set of cards with words relevant to the experience beforehand. During the codesign activity, people sort the cards into groups, or a sequence, or a hierarchy, as you direct them. You might even direct them to give names to the different groups, or steps in the sequence, or layers in a hierarchy. After a few sessions of card sorting, you can analyze those card sorts as a group. The commonalities can be used to build confidence in how to organize the information, whether in a user flow, on a single screen, or during a voice-based conversation. The

differences among the participants' card sorts can illuminate areas that people may find confusing, that require more design effort to include or explain in the overall design.

Another codesign activity is a “magic wand” exercise, in which people are asked, “If you had a magic wand, how would you use it to change this experience?” The power of a magic wand exercise is to uncover the insights about what people actually want to see in an experience. The important part of this experience is not necessarily the new features or interactions they describe, but the conversation that follows. You might learn the problems the person is still trying to solve, and their motivations and concerns about how the problem is solved.

USING SURVEYS

Interviewing 100 people might be too expensive and take too much time, but asking survey questions to 100 people is much less costly. When people respond to surveys, especially to free-response questions, they provide the content designer information about the words they already associate with particular ideas—the words already in their heads.

Many of the questions that can be asked in interviews can also be asked in surveys, though they need to be designed carefully to minimize bias. For example, if a survey offers a fixed set of possible responses to a question, that set needs to be complete for the range of responses the person would choose. For another example, questions that ask for a rating along a scale usually need to have equal weight given to either end of the scale.

Surveying tools include simple forms like SurveyMonkey and Google Forms, or more complex options available through Qualtrics, UserTesting.com, and Optimal Workshop. Specific content questions can help with naming, such as:

- “What would you call a feature that allowed you to...?”
- “Which of these words is the best name for a feature that allows you to...?”
- “If you saw the word *carabiner* on the File menu, what would you expect to find there?”

With surveys, you can learn more about the effect of the text on people's perception of the experience and the brand, in the context of the wider marketplace. In these questions, you might ask people to describe the organization and the organization's competitors. You can compare the words and phrases that they use with the voice concepts that support the organization's product principles and concepts. The closer they match, the more successful the experience is at conveying those principles. When surprises appear, whether positive or negative, it's a good opportunity to analyze the product positioning and the brand attributes to take advantage of the new information.

Another way to gauge people's understanding is to test them, asking them assessment questions at the end of an experience. Assessments are different from surveys, though they can use the same question-answer format. In an assessment, you're trying to find out if the person understands a concept from the experience in the way you intended. This can be extremely important in sensitive areas like privacy or security settings, or in regulated industries like healthcare or advertising, because a test can tell you whether you can expect most people who read the UX text to adequately understand the key concepts. The same tools can be used as for surveys, with additional functions that allow you to designate the correct answer.

Note that writing these questions is an important piece of content design! Good questions (and multiple-choice answers, if you're using those) can't be answered without adequate understanding. That is, they aren't leading questions, giving the person hints about what the correct answer would be, even if the person doesn't really understand the material. If they can accurately and completely answer these well-designed questions about the concepts the UX text is trying to convey, the UX text is adequate for that understanding.

Sometimes, the team wants to learn more but doesn't have time for interviews or codesigning, and wants more specialized information than surveys provide. Frequently, the fastest and least expensive way to analyze the UX content is to apply *heuristic analysis*: to compare the UX content to rules for usability and voice to find the most important places to make improvements.

Apply UX Content Heuristics

When it's not clear where to begin making improvements in the UX content, we can apply general rules about what makes the text in an experience “good.” These general rules, called *heuristics*, can be applied by any expert speaker of the language in which the text is written, when they know the organizational purpose for the experience, the purpose for the people using it, and the organization's voice.

For this book, I've organized heuristics into a generic scorecard that highlights how UX content can be improved. These heuristics draw from my own work and research at Xbox and OfferUp, and in part from Nielsen Norman Group's “10 Usability Heuristics for User Interface Design” (<https://www.nngroup.com/articles/ten-usability-heuristics>). I've used this generic scorecard as a template many times, modifying it as needed for general population and technical experiences.

For this second edition, I've updated the text to add references to the Web Content Accessibility Guidelines (WCAG, version 2.2), which is a set of heuristics designed for website accessibility that many organizations use as a basis for their own accessibility assessments. There are more extensive accessibility requirements based on the WCAG and other standards that don't involve the UX content. Those aren't included in this UX content scorecard because they are more focused on the interaction, visual design, and underlying code. It isn't that those things aren't important, but they aren't the theme of this book.

To use the scorecard, choose a piece of the whole experience that a person would understand as a complete task—for example, finding a route in TAPP, sending a message in The Sturgeon Club, uploading an image to 'appee, or writing a goal in LitMop.

The first thing to do is record the purpose for the experience. This includes both what the person seeks from the experience and what the organization that provides that experience wants. These goals can be in conflict: for example, the organization may want the person to spend as much money as possible, while the person wants to save as much money as possible. Identifying these goals helps us focus on the job that the feature is trying to accomplish. This focus is critical because

it's what the experience will be used for; the usability, then, is how useful the experience is in meeting these specific goals, even when they are in conflict.

Next, you'll systematically use each criterion in the scorecard to examine the experience. This framework gives you a robust way to analyze the experience without getting distracted. For each criterion, record your comments about how well it meets or doesn't meet that criterion. These comments are the most important part of using a heuristic evaluation; they let you demonstrate the application of the entire set of heuristics in a balanced manner and celebrate the best parts while focusing on the ways that the experience could be most improved.

As you go through the criteria, I recommend using a 10-point scale to score the experience for each criterion. That is, if an experience fully meets a criterion, it gets 10 points. If it meets the criterion only a tiny bit, it might get 2 points out of 10. If a criterion doesn't apply, it is left out of the final calculation. This scale is arbitrary: you could choose to use a 3-point scale or a 100-point scale, but I find that a 10-point scale is well understood.

The purpose of the point scale is to enable you to communicate and prioritize the work to be done after the analysis is complete. The scale is subjective, just like there is subjectivity in judging a dog at a dog show. You're using it to record your impressions, not the position on some absolute scale of UX goodness. Knowing that some criteria score a 10 feels good, but it also helps you realize you don't need to focus on those items when looking for improvements. I recommend using the points within the context of a single experience (or a single product, if multiple features within a product are being evaluated) to help you prioritize which parts of an experience to improve.

The scorecard has two main categories: Usability and Voice. Usability criteria are in five groups: accessible, purposeful, concise, conversational, and clear. Voice has the six criteria defined parallel to the voice chart in Chapter 2: concepts, vocabulary, verbosity, grammar, punctuation, and capitalization (Table 7-1). This template and other resources are available online as a Google doc: <https://go.strategicwritingforux.com/ux-text-scorecard>.

TABLE 7-1. A blank template for scoring UX content for usability and voice

UX CONTENT SCORECARD FOR: [FEATURE NAME]			
Person’s goals			
Organization’s goals			
Usability			
Criteria		Comments	Score (0–10)
Accessible	Available in the languages the people using it are proficient in		
	Reading level is below 5th grade (general) or 10th grade (professional)		
	Every element has text for screen readers to speak		
Purposeful	What the person should or can do to meet their goals is clear		
	The organization’s goals are met		
Concise	Buttons have three or fewer words; text is < 50 characters wide, < 4 lines long		
	Information presented or requested is relevant and remembered		
Conversational	The words, phrases, and ideas are familiar to the people using it		
	Information is presented in useful steps, in a logical order		

UX CONTENT SCORECARD FOR: [FEATURE NAME]			
Clear	Actions have unambiguous results		
	How-to and policy info is easy to find and use		
	Error messages help the person move forward or make it clear they can't		
	The same term means the same concept, every time it's used		
Voice			
Criteria		Comments	Score (0–10)
Concepts			
Vocabulary			
Verbosity			
Grammar			
Punctuation			
Capitalization			

To provide an example of using the scorecard, let's examine 'appee's onboarding messages. This is the experience that 'appee provides to help people get started when they first sign up. The messages are displayed the first time someone comes to the main screen.

We'll start by taking screenshots of the experience in the context in which people will commonly encounter it, so that we have a durable copy of what was evaluated.

The three screens of the 'appee onboarding flow start with a message in the center of the screen, over the center of the last challenge-winning image (Figure 7-1). On the second screen, a message points out the bookmark icon in the lower-right corner of that image, used to "save" the image. The third and final screen has a message that points out the "play!" button at the bottom of the screen.

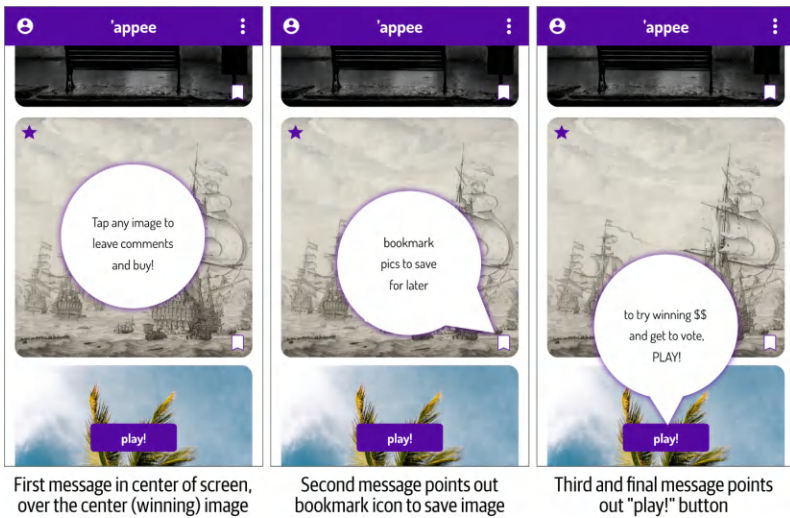


FIGURE 7-1

The three screens of the 'appee onboarding flow are shown in sequence.

To start the scorecard, we need to know what the person is likely trying to do—what is their purpose? And at the same time, what does the organization hope to get out of the experience? Just like for the conversational design exercise in Chapter 3, we begin by listing the goals for the task. For the 'appee onboarding flow, the player's goals are pretty ambiguous: all we know is that they're new to 'appee. The person could be there to upload images, or they could be there to browse images, save them, comment on them, or buy items with images on them.

The business goals for this feature are more straightforward: 'appee wants to reduce the time it takes for a new person to start engaging, to improve the onboarding pace. The key behaviors listed at the beginning of this chapter still apply: 'appee wants people to save, comment, and like images, upload images, browse images, and buy items with those images on them.

These goals fill out the first part of the scorecard template (Table 7-2). This is what the UX content is supposed to be doing for the person and for the organization. To keep the scorecard usable for the team, the goals are kept brief and contextual.

TABLE 7-2. The goals for 'appee onboarding, as entered into the UX content scorecard

UX CONTENT SCORECARD FOR: 'APPEE ONBOARDING	
Person's goals	Ambiguous—could be to save or react to images, play a challenge, start their own profile, or buy an item
Organization's goals	Make sure people know what they can do when they're just getting started, especially saving, commenting, and liking, given that those actions inform the AI and help us personalize ads

Now that we have the screens and we know what the goals are, we can get started scoring the UX content.

ACCESSIBLE

The most essential usability criterion is accessibility. If people can't access the experience, they can't use it! For the purpose of measuring the accessibility of UX content, we have three criteria: language availability, reading level, and labeling.

Available in the languages the people using it are proficient in

Language is the most basic of all forms of access. But even within a region, some significant percentage of a population doesn't speak the dominant language. For example, the US Census reports⁶ that approximately 9% of individuals in the US speak English "not well" and 6% "not at all." If an experience is released only in English in the United States, we might give it a 9.1 on our 10-point scale. This isn't a bad score, but it does acknowledge that the experience isn't usable by at least 9% of the population based on language alone.

Language might seem like an outlier in this feature-specific scorecard, because the language the experience is available in isn't limited to the onboarding messages. The localization languages are typically decided at the product level and outside of the purview of the people working on the UX. But because the language has significant implications for accessibility, it belongs in any assessment of usability.

6 Dietrich, Sandy, and Erik Hernandez. "Language Use in the United States: 2019, American Community Survey Reports." United States Census Bureau, August 2022: <https://www.census.gov/content/dam/Census/library/publications/2022/acs/acs-50.pdf>.

In WCAG 2.2, the criteria 3.1.1: Language of Page and 3.1.2: Language of Parts specifically apply. It's not enough to provide the languages the people need, but they must be able to detect and switch between languages as appropriate.

An additional usability consideration is that localized messages are frequently published more slowly than the native language of the experience. Depending on how the experience is released, there can be a delay of a few days to a few weeks before people will be able to use it in other languages.

For 'appee onboarding, the organization has emphasized language accessibility. They have a hypothesis that people will play more if the challenges are in their native language and their parent corporation operates in several countries. For this criterion, 'appee gets 10 out of 10 points (Table 7-3).

TABLE 7-3. The 'appee onboarding experience gets a perfect score in the language criterion for accessibility in the UX content scorecard

USABILITY	CRITERIA	COMMENTS	SCORE (0–10)
Accessible	Available in the languages the people using it are proficient in	Yes—available in en-US, zh-TW, es-MX, jp-JP, fr-FR, fr-CA ⁷	10

Reading level is below fifth grade (general) or tenth grade (professional)

Reading level is another way to measure the accessibility of language. Even fluent speakers of a language might not be fluent readers of that language. Reading for understanding is a skill that requires practice and more attention than most people pay to the labels, titles, buttons, and descriptions inside an experience. Also, there are differences in cognition across the population that affect a person's ability to read—for example, attention deficit disorders, dyslexia, and concussions. Alcohol- and drug-induced impairment matter, too, if you're making an experience that you want people to be able to use while impaired.

⁷ International standard codes for locales from the International Organization for Standardization (ISO 3166; <https://www.iso.org/obp/ui/#search>) and languages (ISO 639; https://www.loc.gov/standards/iso639-2/php/English_list.php) help us to specify which versions of language the experience is available in.

There are several measures available for reading level in English, including the Flesch–Kincaid Grade Level, Gunning Fog Index, SMOG Index, Automated Readability Index, and the Coleman–Liau Index. These use sentence length and word length to approximate the minimum grade level in the US school system of a person who would understand the text. As far as I know, *none of these measures have been academically validated to use for UX content*. I am eager for colleagues to take on that research and inform us of the results!

Several free calculators exist online for reading level measurement, many of which use a combination of the measures listed. To use some calculators, it helps to add periods to the end of any standalone phrases, buttons, and labels that don’t have them before pasting the modified UX content into the calculators. Sometimes, one of the measures provides very different results than the others, because each of them is calculated differently, but they are often in agreement within half a grade level.

To maximize reading-level accessibility, I keep the reading level below fifth grade for general audiences, and below tenth grade for professional audiences. This is an update from the first edition, based on guidance I learned from practitioners in wider contexts than I had access to in 2019. In general, the lower the reading level, the more usable the content will be—as long as it’s still accurate and complete. For our ‘appee example, the reading level measured by any of the tests is at the second- or third-grade level, well under this limit. For this criterion, the ‘appee onboarding messages get a 10 (Table 7-4).

TABLE 7-4. The ‘appee onboarding experience is scored on the reading level criterion for accessibility in the UX Content scorecard

USABILITY	CRITERIA	COMMENTS	SCORE (0-10)
Accessible	Reading level below 5th grade (general) or 10th grade (professional)	Tests to grade level 2 or 3	10

Every element has text for screen readers to speak

Our final accessibility score is about labeling: each screen or section (WCAG 2.4.2: Page Titled, and 2.4.6: Headings and Labels) and every non-text element on a screen (WCAG 1.1.1: Non-text Content) should have UX text for the screen reader to speak.

This means that any icon, input field, link, or image necessary for understanding should have text that is visible, can be made visible (for example, with hover or mouseover function), or can be made audible (for example, with a screen reader). When information is presented by using visual hierarchy or structure, those visual frameworks also need UX text for the screen reader to speak (WCAG 1.3.1: Info and Relationships).

Along the same lines, the labels that appear to someone visually reading the screen need to be the same as the labels available in the text (WCAG 2.5.3: Label in Name). The various states of each item need to be available, too. For example, it should be possible for a person using a screen reader to know that a button is present but disabled, and to know when that button becomes enabled (WCAG 4.1.2: Name, Role, Value).

The alternate text, whether because of its visible location or audible presence, should be usable to distinguish different actions from one another. In our 'appee onboarding example, the save function is visible in the corner of each image (up to 10 images are newly loaded at a time) on the main screen. But when the experience was tested, the screen reader said “button: Save” 10 times! It's not possible for a person to distinguish which image it might mean, and the experience is reading all the buttons even though they aren't all visible on the screen. This is a bug—a problem that needs to be raised with the engineering team—and it affects the score for this criterion in the UX content scorecard (Table 7-5).

Another impact to this criterion is that the action to take isn't obvious. It's not clear whether the person is being asked to practice taking the action they suggest, tap the message to dismiss it, or tap outside the message to continue. Any of those actions would be common, but the experience needs to specify which it is.

TABLE 7-5. The 'appee onboarding experience is scored on the screen reader criterion for accessibility in the UX content scorecard

USABILITY	CRITERIA	COMMENTS	SCORE (0-10)
Accessible	Every element has text for screen readers to speak	The onboarding message is read, but I can't tell what I'm supposed to do with it. Tap it? Bookmarks just read out "button: Save" 10 times. Can't tell which one is being read, and it's being read all down the screen, even for the ones that aren't visible. Relates to play!, menu, and profile flows.	2

Next, let's decide how well the UX content meets the purpose for this experience.

PURPOSEFUL

Being usable includes meeting the purposes that the organization and the people using 'appee have for this part of the experience. These purposes are recorded in the goals section at the top of the document, but that's not enough. We also need to judge whether the text, as written, will help the person and the organization meet those goals.

What the person should or can do to meet their goals is clear

When the person reads or hears the UX text, it should be clear what the person should or can do to meet their goal. Like we recorded in the last accessibility criterion, the message doesn't make it clear where or whether the person should tap somewhere on the screen. WCAG specifies that "labels or instructions are provided when content requires user input" (WCAG 3.3.2: Labels or Instructions). Without additional indicators, the person doesn't know how to achieve their purpose. In Table 7-6, the comments record the confusion, and it's given a disappointing score of 6/10.

TABLE 7-6. The 'appee onboarding experience is scored on the person's goals criterion in the "purposeful" section of the UX content scorecard

USABILITY	CRITERIA	COMMENTS	SCORE (0-10)
Purposeful	What the person should or can do to meet their goals is clear	Not clear whether the bubble is tappable or whether the person should tap the thing being pointed out. Clear that 'appee wants the person to do something, but not sure how to move forward.	6

The organization's goals are met

The 'appee onboarding experience does a better job of meeting its business purpose than the purpose for the person using it. It points out two of the three specific actions listed in the goals, saving and commenting, but liking is left out.

The text also includes buying, which isn't listed in the goals for this experience. The buying is also out of place because a new player probably wouldn't understand what they would be buying. In Table 7-7, the scorer's comments record both where the text is aligned to the organization's goals and where it isn't.

TABLE 7-7. The 'appee onboarding experience is scored on the organization's goals criterion in the "purposeful" section of the UX content scorecard

USABILITY	CRITERIA	COMMENTS	SCORE (0-10)
Purposeful	The organization's goals are met	The app points out to tap images to leave comments and buy, but not what they'll be buying; bookmarks cover saving, not liking.	8

In UX content, it's not only important that the text is accessible and meets the purpose. There is limited space and attention for the words, and nobody came there to read them. To be usable, the text must also be concise.

CONCISE

There are two ways to measure concision in the UX content scorecard: the overall length of the text and whether the text includes ideas that aren't relevant to the person reading it.

Buttons have three or fewer words; text is fewer than 50 characters wide, fewer than four lines long

Whether the text is on a large screen (like a television) or a small one (like a mobile phone), brief messages are easier to scan. The text that has tested well, in my experience and in proprietary research that I've consumed, has been three or fewer lines long and no more than 50 characters wide. In the same vein, buttons that have one or two words are used more often and faster than buttons with more words. These are challenging limits to meet, but well worth the effort. The 'appee onboarding messages meet all three of these subcriteria as currently written, without even needing a comment (Table 7-8).

Information presented or requested is relevant and remembered

Limiting the information to just what is relevant to the person may be the most challenging part of writing UX. There are two kinds of difficulties: first, sometimes we know several different ideas that would be relevant, depending on who the person is. But because we can't always make different experiences for different people, or even know which person is which, we usually need to write one piece of content that every person will encounter. Second, sometimes it's not possible to know what the person would consider relevant. In either case, we just have to do the best that we can.

When we're requesting information, we have a similar constraint: only asking for the information we need and only asking for it once (WCAG 2.2:3.3.7: Redundant Entry). We need to understand why the experience is asking for any piece of information: which purposes does it serve for the person and for the organization?

For the 'appee onboarding experience scorecard, we can note in comments that we can't be sure what the person actually wants, but we also don't know that the ideas we are presenting are relevant to them (Table 7-8).

TABLE 7-8. The 'appee onboarding experience is scored on the two criteria of the concision section of the UX content scorecard

USABILITY	CRITERIA	COMMENTS	SCORE (0-10)
Concise	Buttons have 3 or fewer words; text is < 50 characters wide, fewer than 4 lines long		10
	Information presented or requested is relevant and remembered	We're not sure what they want. But because they're new, they should at least be introduced to what's possible. Getting to vote feels out of place. Liking is left out.	8

If an experience were fully meeting its purposes and fully concise and nothing else, it runs the risk of being robotic. That roboticness, by itself, can make the experience more difficult to use, which is why the text in the experience also needs to be conversational.

CONVERSATIONAL

There are many aspects of making an experience feel conversational that have to do with voice, which has its own heuristics. For the parts of conversation that are about usability, it's most important that the experience is using words and concepts that the person would be familiar with, and that those ideas are brought up in a sequence that makes sense.

The words, phrases, and ideas are familiar to the people using it

When an experience uses the words that are most familiar to the person using the experience, the usability of the experience skyrockets. Separate from the terminology (which is in the “clear” section, next), the day-to-day words, phrases, and grammar that people would use to describe the experience to their friend or family member are the words that will be most understandable. It's also possible to provide definitions for unusual words, idioms, jargon, and abbreviations (WCAG 3.1.3: Unusual Words and 3.1.4: Abbreviations) so that people new to the experience or intellectual domain will be able to understand it.

In our 'appee example, most of the ideas brought up by the onboarding messages should be familiar to the person if they have used social media or social games in the past. This is an expected background for a person to have before using 'appee for the first time. The one potentially unusual idea is voting: people voting for their favorite challenge image is not a common idea, so a person new to 'appee is not likely to have encountered it before. Bringing it up in the onboarding experience, before that person has the context of uploading images, makes for a weird moment in an otherwise smooth conversation. For this, 'appee loses one point for this criterion (Table 7-9).

TABLE 7-9. The 'appee onboarding experience is scored on the familiarity criterion in the conversational section of the UX content scorecard

USABILITY	CRITERIA	COMMENTS	SCORE (0-10)
Conversational	The words, phrases, and ideas are familiar to the people using it	Voting might not be totally familiar here, as an idea	9

Information is presented in useful steps, in a logical order

It's important to not only use the words that people easily understand, but also to present the ideas to them in the most helpful, logical order possible (WCAG 1.3.2: Meaningful Sequence). People are significantly more successful when ideas are put in the order in which they need to use them. For example, consider the difference between these two statements:

- “To allow location, go to Settings, then turn on Location.”
- “To allow location, turn on Location in Settings.”

The second sentence is shorter but puts information in the wrong order. The sequence the person will need to find them is first Settings and then Location.

The 'appee onboarding messages are probably not in the correct order. Even though these directions are not intended to be followed immediately, they neither build on one another nor mimic a person's likely path through the experience. As the comment in the UX content scorecard

suggests, it would make more sense to start people with the interactions that are easiest and take the least commitment, before diving into actions that require a higher level of commitment (Table 7-10).

TABLE 7-10. The 'appee onboarding experience is scored on the logical order criterion in the conversational section of the UX content scorecard

USABILITY	CRITERIA	COMMENTS	SCORE (0-10)
Conversational	Information is presented in useful steps, in a logical order	People are super unlikely to buy as their first action or to leave comments. Order should probably be Save > Like > Comment? Buy might not be the core action...	4

As much as usability is created by creating text that is purposeful, concise, or conversational, it's twice as important for that text to be clear. Without clarity, people can feel good and follow the path but still not have confidence that they understand it.

CLEAR

Clarity is a metaphor: it literally means transparent, like clean glass or an unobstructed view. Applying that metaphor to the UX content, it means that the words are doing all that they can to help the person understand the experience. The person has the information they need, and that information makes sense so that they can meet their goal.

The four criteria in the “clear” section check for symptoms of systematic clarity. These are attributes that should be present in each part of an experience, from sign-up to the “Sad to see you go” screen. When they are missing in one part, the whole experience suffers. But even though they are systematic criteria, they can apply more or less in each part of the experience. They can also be fixed part by part, message by message, until the whole experience is clear.

Actions have unambiguous results

Most experiences need to use the UX text on the buttons, titles, controls, and more, so that the person expects the actions that happen. For example, if a person is checking a checkbox, there is a visible change, and a screen reader speaks “checked.” People should also see or hear when the action is complete. In WCAG 2.4.4: Link Purpose (In Context), hyperlink text is noted specifically: “The purpose of each link can be determined from the link text alone...”

When someone takes a more substantial action, it can be important to give them more robust feedback. For example, when a “Pay Now” button is available in an online purchase, they can reasonably expect that money will be taken from their account. After they use it, the person expects some sort of confirmation that the purchase is complete.

Now, imagine that the button at the end of the purchase flow said “Next” or “Continue,” and the next thing the person saw was a confirmation that the purchase was complete. The “Next” action had an ambiguous or even misleading result because the person wasn’t reasonably expecting to commit to the purchase at that moment.

In the ‘appee onboarding messages, people can’t determine whether they should tap the message or not. The text inside the messages gives a little more clarity to what will happen when the bookmark icon is tapped, but the message makes it seem as if the act of tapping an image will either leave a comment or buy an item with that image on it. Even the relatively clear “play!” button becomes less clear, because the message that points to it conflates winning money and voting. Of the 10 points possible, ‘appee onboarding messages get only 2 points (Table 7-11).

TABLE 7-11. The ‘appee onboarding experience is scored on the ambiguity criterion in the “clear” section of the UX content scorecard

USABILITY	CRITERIA	COMMENTS	SCORE (0-10)
Clear	Actions have unambiguous results	Not clear what people should do here, but save text seems clear. PLAY is less clear, because of both winning and voting. Tapping the image sounds like it might leave a comment or buy something?	2

How-to and policy information is easy to find and use

In any population of people who will use a software experience, there will be people who are comfortable tinkering with the experience until they find what works, and there will be people who don't want to click (or tap) anything until they're sure they will do it right. Most of the people who build software, and those who are in charge of most of the decisions about how an experience is made, are in the first group: we are tinkerers.

Tinkerers like us tend to think that if we're doing our jobs right, people shouldn't need additional help. We take it as a badge of honor when we can set up something new without using the instructions, whether it's an electronic gizmo or an IKEA table. An experience that needs to include how-to information is broken, our thinking goes, because the experience should make it clear what the person should or could do.

But quite a bit of the population aren't tinkerers, according to Margaret Burnett's work on GenderMag (<http://gendermag.org>).⁸ Most of the people who could be using software aren't the tinkerers who make software. Most people want to understand the experience before they are comfortable clicking, tapping, or experimenting.

We tinkerers need to get out of our own biases enough to include these people (and their spending power!) in the experiences we create. To increase usability for all audiences, we can make sure the how-to and policy information is easy to find and use. When we do this, we not only make more people capable of using the experience, but we also make it easier for more people to enjoy it.

The 'appee onboarding messages are themselves how-to information. There's also help information available in the menu, under Help. On the UX content scorecard, 'appee onboarding messages score 10 out of 10 points for their findable how-to information (Table 7-12).

8 Mihaela Vorvoreanu et al., "From Gender Biases to Gender-Inclusive Design: An Empirical Investigation," *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (April 2019), <https://web.engr.oregonstate.edu/~burnett/Reprints/chi19-GenderMag-findToFix.pdf>; Margaret Burnett, Anicia Peters, Charles Hill, and Noha Elarief, "Finding Gender-Inclusiveness Software Issues with GenderMag: A Field Investigation," *CHI '16: Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (May 2016): 2586–2598, <https://web.engr.oregonstate.edu/~burnett/Reprints/chi16-GenderMag-fieldStudy.pdf>.

TABLE 7-12. The 'appee onboarding experience is scored on the “availability of help” criterion in the clear section of the UX content scorecard

USABILITY	CRITERIA	COMMENTS	SCORE (0-10)
Clear	How-to and policy information is easy to find	This is how-to info; no policy needed right now	10

Independent of the divide between tinkerers and nontinkerers, a common place that many people want help is when they reach error conditions. But even better is when the error message is so clear that additional help isn't needed.

Error messages help them move forward or make it clear that they can't

When a person hits the end or edge of an experience, the experience usually displays an error message (WCAG 3.3.1, Error Identification). This can be as usable and clear as telling them what to do, or as confounding as describing an underlying technical condition the person has no knowledge of, and no way to affect.

As described in the Errors section in Chapter 4, error conditions are among the most important places to empathize with the person. The person is trying to use the experience. They may be there for entertainment, for work, for civic responsibility, or as a chore. Whatever they are there for, the error is preventing their progress. The kindest and most usable thing the experience can do is to help the person move forward to meet their goal, despite the error (WCAG 3.3: Error Suggestion). If they can't be moved forward, the error message should maximize usability to make that clear, so the person can find a different way to meet their own needs.

Just like the other criteria, if the 'appee onboarding messages included the possibility of errors, we'd score them as a whole in the UX content scorecard. For example, if there were 10 possible errors, and only 8 met the criterion, the experience would get 8 points. Because the 'appee onboarding messages don't have any error conditions to use as an example, this criterion is marked “not applicable” (N/A), and those points don't count toward the total (Table 7-13).

The same term means the same concept, every time it’s used

Our last criterion for clarity is the terminology. Different from other parts of the vocabulary, a *term* is a word set aside by the experience as having a specific meaning. Terms need special treatment in UX content so that the same idea is always called the same term, and the term isn’t used for other concepts even when those concepts are similar in the experience (WCAG 3.2.4, Consistent Identification).

In ‘appee, the term for saving an image is “bookmark.” Because the onboarding message correctly uses “bookmarks” instead of “saves” (and that’s the only term that appears in the messages), the ‘appee onboarding messages get points for being consistent about the way it uses the term (Table 7-13).

TABLE 7-13. The ‘appee onboarding experience is scored on the last two criteria in the clear section of the UX content scorecard

USABILITY	CRITERIA	COMMENTS	SCORE (0-10)
Clear	Error messages help the person move forward or make it clear that they can't	No error conditions in this flow	N/A
	The same term means the same concept, every time it's used	Bookmark vs. save	10

Although we are done with the usability score, there’s another important lens to use: voice.

VOICE

Usability is about two-thirds of the possible scores of the UX content scorecard, and voice is the other third. This corresponds well to how we think about goals. The organization and the person using it need the experience to be usable, but the organization is the only party that meets some of its goals through using a recognizable voice. The person benefits from the recognizable voice, but it is unlikely to help them meet their goals.

The criteria for measuring voice come directly from the organization’s voice chart (Chapter 2), aligned to the concepts, vocabulary, verbosity, grammar, punctuation, and capitalization it defines. If the organization doesn’t have a voice chart, it would include whatever brand criteria it does have, so that the experience can be analyzed through this important lens.

We defined different aspects of the voice to align with different product principles, so we need to choose which product principles apply to this part of the experience. For example, 'appee has three product principles: Playful, Insightful, and Surprising. The Insightful product principle doesn't apply here, because 'appee relates that principle only to the images that people upload and comment on. Therefore, the scorecard for the onboarding messages need to include only the voice attributes that apply exclusively to Playful and Surprising in the Criteria column (Table 7-14).

The same way that we did for the usability section, we can score the 'appee onboarding experience for voice. Because the Insightful principle was removed, there's no specific guidance for vocabulary, so those points are N/A (not applicable). Comments indicate that a few points are lost because sentences were used instead of phrases, a word doesn't need capitalization, and more emojis could be used throughout (Table 7-14).

TABLE 7-14. The “voice” section of the UX content scorecard for the 'appee onboarding messages

VOICE	CRITERIA	COMMENTS	SCORE (0-10)
Concepts	Small delights, avoiding grand successes; frippery Unpredictable; misdirection and difficulty can be fun	Lacks small delights, no frippery. There's difficulty, but is it fun?	2
Vocabulary	{Not vocabulary-specific}		N/A
Verbosity	Fewer than strictly necessary	These are brief, but they don't leave me guessing in a good way	8
Grammar	Present and future tense Phrases preferred	Could be phrases instead of sentences	8
Punctuation	Avoid periods; use emojis, exclamations, interrobangs, question marks	OK, but why not more emojis?	9
Capitalization	Use capitalization only for emphasis	Inconsistent—does “Tap” need to be capitalized?	9

The most interesting criteria to discuss here are the concepts, a category for which these 'appee onboarding messages get a dismal score. There's no information or small delights included in the messages. There is some difficulty, but the difficulty feels more like usability problems than an intentional challenge. Even though, from a usability perspective, the messages could have been clearer, the concept criteria would make them even more challenging. This is one of many places where, for 'appee, the voice and usability are at odds.

This tension is natural. Throughout human experiences, there are design criteria that contradict one another. Consider them as ends of a spectrum within which the voice of the experience exists. At different points in the experience, different points on that spectrum of voice are appropriate. Examples are all around us: traffic signs are both highly visible and not distracting. Museums collect and preserve artifacts but also display and use them. Hospital devices can attract the attention of medical staff but let patients sleep.

Games are a special case in which usability is intentionally thwarted. Part of what makes a game fun, whether it's a puzzle or a first-person shooting game, is that there is inherent challenge. That challenge isn't always reflected in the words, but 'appee is constructed to make the words challenging. By noting the balance between voice and usability scores, a team can record where and how those decisions were made and choose to make adjustments.

We've now completely scored the 'appee onboarding messages with our UX content scorecard. We can add up the 125 points earned and divide them by the 170 points possible for this experience to get the total score of 73% (Table 7-15).

TABLE 7-15. The completed UX content scorecard gives the 'appee onboarding experience a total score of 73%

UX CONTENT SCORECARD FOR: 'APPEE ONBOARDING			
Person's goals		Ambiguous—could be to save or react to images, play a challenge, start their own profile, or buy an item	
Organization's goals		Make sure people know what they can do when they're just getting started, especially saving, commenting, and liking, given that those actions inform the machine learning model and help us to personalize ads	
Usability			
Criteria		Comments	Score (0-10)
Accessible	Available in the languages the people using it are proficient in	Yes—available in en-US, zh-TW, es-MX, jp-JP, fr-FR, fr-CA	10
	Reading level is below 5th grade (general) or 10th grade (professional)	Tests to grade level 2 or 3	10
	Every element has text for screen readers to speak	The onboarding message is read, but I can't tell what I'm supposed to do with it. Tap it? Bookmarks just read out “button: Bookmark” 10 times. Can't tell which one is being read, and it's being read all down the screen, even for the ones that aren't visible. play!, menu, and profile work.	2

UX CONTENT SCORECARD FOR: 'APPEE ONBOARDING

Purposeful	What the person should or can do to meet their goals is clear	Not clear whether the bubble is tappable or whether the person should tap the thing being pointed out. Clear that 'appee wants the person to do something, but not sure how to move forward.	6
	The organization's goals are met	It points out to tap images to leave comments and buy, but not what they'll be buying; bookmarks cover saving, not liking.	8
Concise	Buttons have three or fewer words; text is < 50 characters wide, < 4 lines long		10
	Information presented or requested is relevant and remembered	We're not sure what they want. But because they're new, they should at least be introduced to what's possible. Getting to vote feels out of place. Liking is left out.	8
Conversational	The words, phrases, and ideas are familiar to the people using it	Voting might not be totally familiar here, as an idea	9
	Information is presented in useful steps, in a logical order	People are super unlikely to buy as their first action. Nor to leave comments. Order should probably be Save > Like > Comment? Buy might not be the core action...	4

UX CONTENT SCORECARD FOR: 'APPEE ONBOARDING'			
Clear	Actions have unambiguous results	Not clear what people should do here, but bookmark text seems clear. PLAY is less clear, because of both winning and voting. Tapping the image sounds like it might leave a comment or buy it?	2
	How-to and policy info is easy to find and use	This is how-to info; no policy needed right now	10
	Error messages help the person move forward or make it clear they can't	No error conditions in this flow	N/A
	The same term means the same concept, every time it's used	Bookmark vs. Save	10
Voice			
Criteria		Comments	Score (0-10)
Concepts	Small delights, avoiding grand successes; frippery Unpredictable; misdirection and difficulty can be fun	Lacks small delights, no frippery. There's difficulty, but is it fun?	2
Vocabulary	{Not vocabulary specific}		N/A
Verbosity	Fewer than strictly necessary	These are brief, but they don't leave me guessing in a good way	8
Syntax	Present and future tense Phrases preferred	Could be phrases instead of sentences	8

UX CONTENT SCORECARD FOR: 'APPEE ONBOARDING			
Punctuation	Avoid periods; use emojis, exclamations, interrobangs, question marks	OK, but why not more emojis?	9
Capitalization	Use capitalization only for emphasis	Inconsistent—does “Tap” need to be capitalized?	9
Total points			125
Points possible			170
Score			73%

This raises the question: is 73% a good score? The scorecard is a proxy measurement for how well the UX content meets the goals of the organization and the people who use the experience. The score lets us predict how much we can improve the quality of the text, by putting a value on how closely the UX content comes to being usable and in the voice defined for the experience.

More important than the numeric score, we have performed an analysis that has identified steps that we can take to improve the UX content. Our hypothesis is that we can take these steps to improve UX content quality, which will improve the experience’s ability to help the organization and the people meet their goals.

Here are the lowest scores from these criteria:

- **Accessible:** Every element has text for screen readers to speak.
- **Conversational:** Information is presented in useful steps, in a logical order.
- **Clear:** Actions have unambiguous results.
- **Concepts:** Small delights, avoiding grand successes; frippery; unpredictable; misdirection and difficulty can be fun.

The team, armed with this information, can make decisions and prioritize the work to be done to improve the experience. They can score other parts of the experience, like playing a challenge, buying an item, or leaving a comment, and decide which parts need the most investment to achieve a higher score.

As the team makes improvements, it would be wise to conduct research so that the impact of the changes can be identified (and in some cases, quantified.) When 'appee measures its improvements in engagement, retention, cost reduction, and other direct measurements, the team can identify how making UX text changes, based on their heuristic analysis, helps them drive the business outcomes they seek.

Summary: If You Like It, Put a Value on It

You may have heard the quotation (or seen my second-hand misquote of it in the first edition): “If you can’t measure it, you can’t manage it.” The full quote by W. Edwards Deming is “It is wrong to suppose that if you can’t measure it, you can’t manage it—a costly myth.”

The truth I want to convey here is that we can’t make improvements when we don’t have feedback. We can make changes, but we can’t determine whether those changes are good or bad if we aren’t paying attention. This chapter has been all about methods of paying attention, whether as quantitative measurements or qualitative judgments.

As we make improvements to the UX content, we should realize some improvements, for example, in engagement, completion, retention, referrals, and speed to onboarding. We may also realize improvements in brand affinity, perception of quality, and other positive effects that are more difficult to quantify. Whether measurable or not, each improvement is worthy—and they add up.

We can get closer to understanding why UX content is effective when we conduct research, including interviewing people and analyzing their complaints and questions. People will tell us their feelings, their preferences, their likes and dislikes. They’ll tell us how the experience works for them, and they’ll usually be as accurate as they can. But just like the people on the team, the people who use the experience can be wrong about why they do what they do, and why they like what they like.

That’s why the heuristic measures are valuable, even when the team can also use A/B testing and research. The usability heuristics are a set of guidelines that are generally true about UX content, regardless of an individual person’s beliefs about why they behave in certain ways. The voice heuristics are a set of guidelines about what the organization believes is true for them, their experience, and the people who use their experience. Together, the heuristics are a hypothesis of why and how

the UX content can be good. When they are applied to an experience, the scorecard gives us a path forward to fixing the words—and helps us to understand when it's not the words that need to be fixed.

Tools and Processes

It's best to have your tools with you. If you don't, you're apt to find something you didn't expect and get discouraged.

—STEPHEN KING, WRITER

THERE ISN'T ANY ONE magic tool or rigid process to be a good content designer. Most content designers use widely available tools, even tools that are available for free. We can define our own processes, and we need to adapt and align to our teammates' and organization's processes. In this chapter, I share the tools and processes I use to be successful while I accomplish the main tasks of content design:

- Write for the context
- Manage content review
- Publish the text
- Track the content work to be done
- Systematize content design

Write for the Context

This book has sentences grouped in paragraphs, organized in sections and subsections, to do its word-based work. I used two common word-processing programs to write it in approximately the same format you're reading it now.

In contrast, content inside a user interface is not a sequence of words, sentences, and paragraphs that stand on their own. Instead, it exists to be the conversation between the experience and the person using it, living in the labels, headings, descriptions, and more. The experience talks to the person with words, information hierarchy, and visuals, and the person responds by interacting with elements on the screen.

To choose the appropriate words while writing for interfaces, we need to consider not only titles, sections, and paragraphs, but also buttons, controls, flyouts, dialogs, text input fields, and more. Our words can be seen, heard, or both. When the person encounters this writing, they don't progress from the top to the bottom of a screen, but rather their eyes flick from top to bottom, from title to button, and might skip any of the words we write.

So, why would we imagine that we could use tools designed for document writing to create content for a digital experience? If I were to write the UX text in a written-page format, or even in tables or a spreadsheet, I would fail to design for the context.

Instead, I need to set myself up for success: I need to write and evaluate the content where people will encounter it. Then, I will need to try many options and save many iterations.

WORKING IN DESIGN TOOLS: FIGMA, SKETCH, AND SIMILAR TOOLS

When I'm working with a designer, whether it's on a new experience or updates to an existing one, they usually have a working file in a UX or graphic design tool like Figma, Proto.io, or Photoshop. When I work in the same tool, I can iterate the UX text more rapidly and communicate with designers more effectively than any other way.

In an ideal sense, it is great to know what's being designed, and be able to work alongside the designer while the UX is being designed. Tools like Figma make that possible, with shared views of the same designs at the same time. Even before Figma, though, the essential ingredient wasn't the tool but the commitment a designer and writer make to the collaboration. Some designers and content designers can pair-design the way some engineers pair-program. But even when they can't or don't work simultaneously, the designs and UX text can be revised by both designer and writer, iterating the text and designs in turn.

In many ways, using a design tool is similar to using text boxes to create graphics or slides with text. It doesn't matter whether the text is standing alone or part of a group, symbol, or component. I just need to be able to get to the text box (or insert a text box) and edit it. Sometimes, I need to mock up a new screen to adequately convey how the text should work. Design tool knowledge is essential for content designers.

When I'm working in a design tool, I take advantage of the (almost) infinite canvas provided. When I do the content design work, I work iteratively. This means that instead of taking one instance of a message and polishing it, whittling it down, and editing it in place, I make a copy of the whole screen in which it appears and change it. Then I make a copy of the new screen and change that one (find examples of this iteration in Chapter 6).

By working in this way, I create several possibilities for the content. It gives me more freedom to take bigger chances and make bigger differences—and think deeply about the content and how it will serve its purposes. I can annotate the options in the design file, recording what I'm trying to do: leverage the voice, emphasize different ideas, or make sure I convey the right information.

Eventually, I save and annotate my best options and share those with the designer and the rest of the team. I'm likely to share them where the team is working, whether that's in slide decks, the design file, in messaging, or elsewhere. I also have a set of terrible and mediocre options that I can keep in my back pocket so I can show the various ideas I've tried.

LO-FI TOOLS: PRESENTATION, PHOTO EDITING, AND SIMILAR TOOLS

Frequently, the content designer needs to work on words that have no design file prepared. Maybe that is because the screen was built from a common framework, so developers didn't need pixel-perfect images. Or maybe it was built long ago or by a different team. For any number of reasons, you don't have a design to edit. You can use a screenshot or build up your own screen in the design software, or if that isn't available, you can do the same in lower-fidelity tools.

Consider this scenario: you were sent a screenshot, attached to a bug report, email, or Slack conversation, that says, "This doesn't sound right to me..." For example, Figure 8-1 is a screenshot of a poor error message in 'appee, illustrating a common problem: errors are often written from the perspective of the engineer instead of being designed to be understood by the person using the experience.

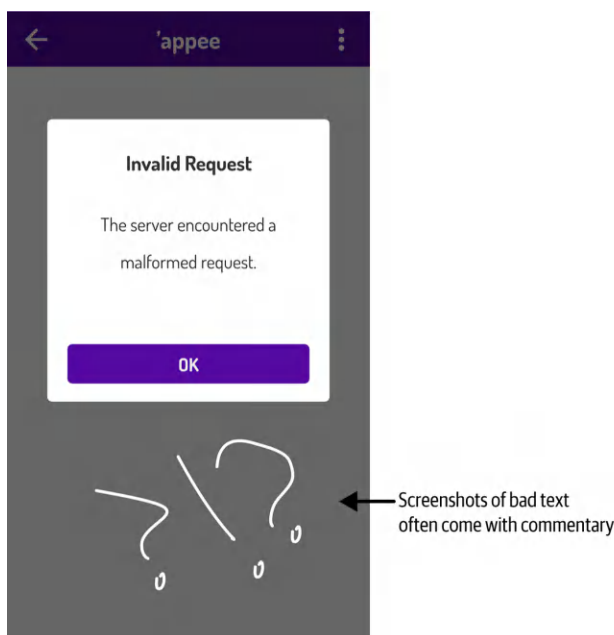


FIGURE 8-1

This 'appee screenshot has an error message that's incomprehensible to the player.

Unfortunately, words captured in a screenshot aren't editable: they exist as pixels, not as text that can be typed, deleted, and edited.

To make a version that I can edit, I begin by pulling the image into any of these pieces of software: Figma, Canva, Keynote, Microsoft PowerPoint, or Google Slides, or even Paint. Any of these will work because they all have these two tools available: text boxes and rectangles.

My goal isn't to make the new text pixel-perfect. To get the text right, I only need the pixels to look good enough that they don't distract from evaluating the new text options that I write. I can do almost the same work using pencil and paper (and do, if I need to!), but the electronic version will be more useful to iterate, share, and convince others.

In any of these pieces of software, I can create an editable version by drawing a text box over the text that I want to change. It might seem backward, but I begin by typing in the same text that already exists. Then, I adjust the font, size, and style of the text until it matches the original. This gives me the constraints of the text width and depth

available in the current design so that I know how much text I could fit there. Sometimes I don't have the correct font installed, so I choose one that is close. Again, my goal isn't perfection, but to get close enough to evaluate.

This text box will be the top layer of a three-layer image. The screenshot itself is the base layer. Now I need the middle layer to block out the existing text, so I draw a rectangle the same size as the original UX text and make it the same color as the background color of the text. I now have three layers: the text on top, then the plain rectangle that blocks out the old text, and then the original image underneath (Figure 8-2).

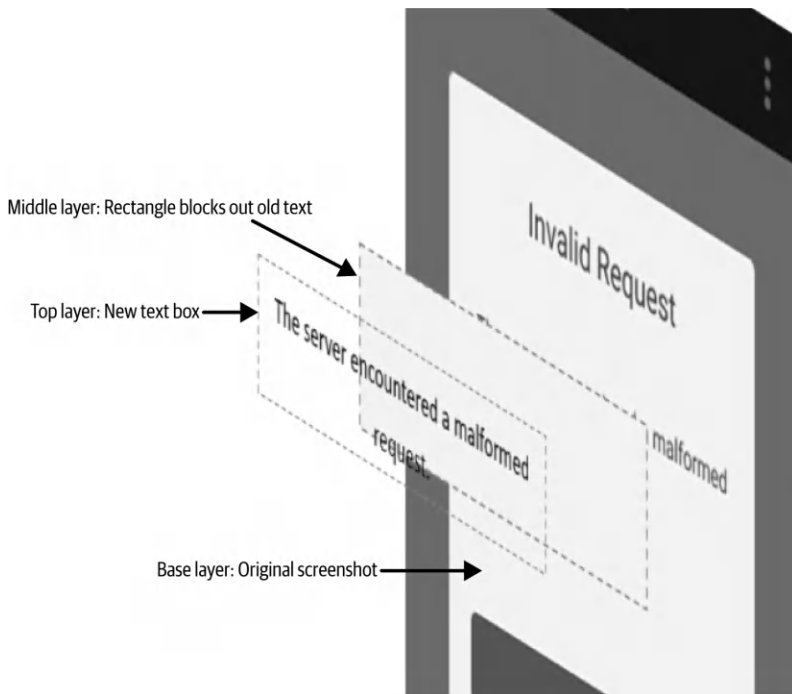


FIGURE 8-2

An editable text box and a rectangle are layered on top of an original screenshot.

To write my drafts, I make a copy of the entire group and then edit the text (Figure 8-3). Then, I make another copy and edit the text again, following the iterative process I described earlier. In general, I'll edit and make copies until I have several good options to choose from.

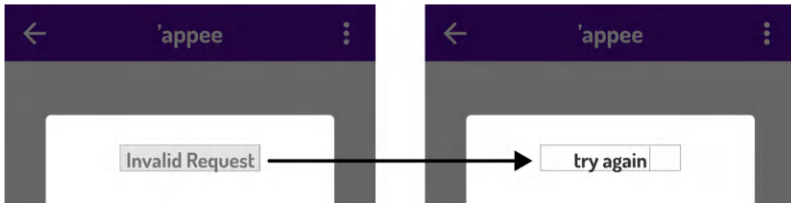


FIGURE 8-3

To start editing, I try out new words in the text boxes I just added to the screenshot.

I continue to iterate new options until I have a few that I think will work (Figure 8-4). For more detail about the editing process, go to Chapter 6.

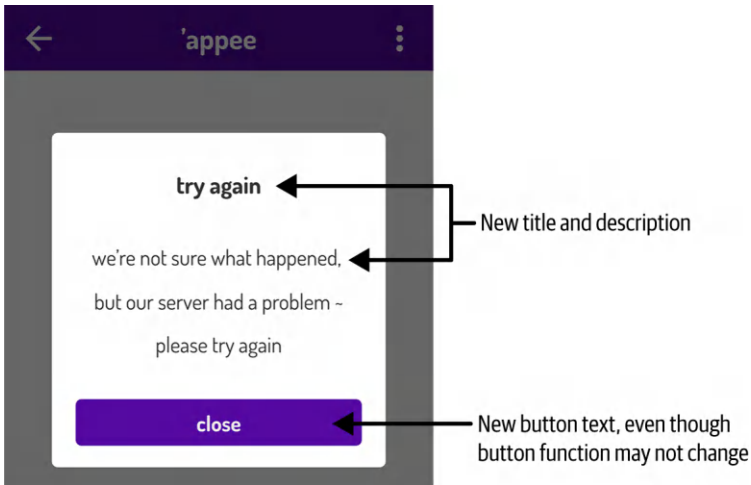


FIGURE 8-4

The result of editing the error message in Figure 8-1 includes changes to every piece of text: title, button, and description.

I take the best of the options and share them with the team in a content review document (described in the Manage Content Review section, coming up). Allowing the team to review the UX text in the context of the design makes it possible for them to understand the impact the new words will make.

USING AI TO WRITE UI TEXT

Chances are, you've been using LLMs as writing assistants for a long time. Early versions of spell checking and grammar checking used lists; if a word was on the "common misspelling" list, you could set it to be autocorrected for you. Now, we have LLM tools embedded into most authoring tools. They scan words and grammar to suggest more usual alternatives. As I write this book, common typos are corrected before I even notice them, and blue squiggles underline content that may be incorrect.

One problem is that these automatic corrections and suggestions are sometimes incorrect. Sometimes words get automatically changed to be the wrong word entirely. Many of the suggestions are just plain wrong—even when writing content in sentences and paragraphs, which is the most common data type for these LLMs to be trained with. The person designing the content needs to understand the nuances of different words for the people who will read them.

LLMs can definitely generate content faster than humans. However, the LLM can't decide that it's the right content, nor can it imagine the context of the person who will read it. As discussed in Chapter 5, the LLMs use statistics to guess whether the text follows the same patterns it finds most probable. When we're solving problems in new features or in technical spaces, or leveraging the product's unique voice, we need what's most effective, whether or not it's probable.

To elicit the right content from LLMs, the content designer not only needs to understand the problems to be solved and how people will use the experience, but also has to formulate those understandings into robust prompts for the LLM. Then, the LLM can provide multiple options for the content designer to think through, consider, and work to refine. Finally, the content designer needs to come up with explanations for those options to share with the team, plus reasons why they'd recommend one over the other, even though they have missed the thinking process of creating them in the first place.

In contrast, the content designer could start from text patterns like those in Chapter 4, which can be tailored to include in product design systems. Instead of needing to translate their knowledge about the context and the purposes into prompts, the content designer could spend

that creative time, well, creating. By recording their explanations alongside their designs, they have everything they need to persuade their team, including the rationale behind the recommendations they make.

Writing can be described as the process of thinking while recording those thoughts. When designing content for user experiences, the thinking is a much slower part than the writing! If you are using an LLM to create UX content, make sure that you're reserving time and energy to think through the problems and solutions, even while the LLM generates the sequences of words.

Manage Content Review

After the UX content is drafted, it generally needs to be reviewed by a wide range of teammates. This could include the engineers, UX researchers, designers, product managers, attorneys, marketers, and more.

This larger process is not cocreation with the wider team, because then the content would suffer from all of the common “written by committee” problems. For example, the text can become so generic that it becomes meaningless, or it can get stuffed full of so many competing details that people won't read it.

Instead, content review is the process of collecting and addressing the feedback, ideas, and concerns of each team member. It's important to believe that the entire team is trying to make the best possible experience for the people who will use it, and drive the best outcomes for the organization. The person designing the content needs to explain, persuade, negotiate, and make deliberate improvements judiciously.

Whenever possible, I present the content to be reviewed in a meeting. This can be in an existing design meeting, feature team meeting, formal critique session, or a separate meeting that I schedule for the appropriate people. Like any creative work, I'll get a more effective review of the content when I help people understand the choices behind it. My work can stand on its own, but it can't tell the whole story, including its context and the problems it needs to solve. It also can't respond with answers to questions or incorporate real-time changes—but I can.

Sometimes, the design work has been so collaborative already, a separate presentation isn't needed. I will still record the rationale for my recommendations, including the problems to be solved, relevant research,

and more, usually as text-box annotations to the designs (not in comments). Then, my teams and I use the comments to ask questions and make specific suggestions.

Comments can quickly get overwhelming as people reply to them, creating long threads of comments and replies. Before I present the design, I explicitly establish how I'll use comments. For example, I ask people to use separate comments for separate ideas. I let them know that as those questions or problems are resolved, I will close the comment. This helps me clean up the work so that we know which problems still need resolution and which have been fixed.

The most challenging feedback, but sometimes the most common, comes during unrelated meetings or in messages (chat, email, etc.). This feedback is still critical to address! But because those criticisms can come from anywhere, it's important to record them with the other feedback, as comments in the design.

Document Regulated Content

In some cases, there's a business or regulatory requirement to document the content that exists throughout the experience. In my work, this has most often appeared in health-related technology, though I've also seen it in privacy and security work. In these cases, we create text-based review documents that need to be created and maintained separately from the designs.

The purpose of the review document is to create a bridge between the smaller team that collaborates closely on the design and a larger group of reviewers, including potential legal and regulatory reviewers, outside of the organization. The review document allows team members to not only make suggestions and comments, but also discuss those suggestions and comments in one place, asynchronously. Those comments and suggestions become part of the document, in its digital history. At the end of the design process, the text from the document can be copied and pasted (instead of retyped) into code.

The tools that I've found to be easiest and cheapest to use are online versions of Microsoft Word or Google Docs (or similar). This relatively manual method of storing and sharing documents creates an online record so that asynchronous comments can be shared and seen by the entire review team.

In the review document, I create a table. On one side is an image of the screen. Beside it, I list the text for the screen in an editable and commentable form. I use separate rows for different elements of the design so that titles, descriptions, buttons, labels, or other elements have their own cells within the table.

For example, Figure 8-5 has the content review document for April challenges in 'appee. Above the table, I put the context for the document—in this case, how many challenges are needed and the overall theme for the month. (If it were a more complex project, I might also include links to the design files or other product docs.) Then, I copy and paste images and text from the design document.

Challenges for April 2019

April summary:

- 30 challenges with 3 emergency backups
- Texture and color; make the most of springtime colors, Easter; note that it's autumn below the equator!

Design	Element	Suggested text
1.	Title	FUZZY
	Description	floofy or feathery, blurry or baffling What makes you feel warm or cozy, or dazed and confused?
2.	Title	fresh as a
	Description	breeze, baby, daisy, clean laundry, radish, car wash, new guy at work

FIGURE 8-5 April challenges are organized in a content review document for the 'appee team to review.

When the document is ready, I send it to all of the reviewers. Usually, this is in email or on a group messaging system, where I give the link and the timeline for review. For example, I might write an email to the reviewers:

Subject: Review April challenges before noon, March 25

April challenges are ready for your review!

These need to be coded before March 26 to be tested. If I don't hear from you before **12 noon on March 25**, I'll take that silence as approval.

Please comment in this doc: {link}

When people comment and make suggestions, everyone on the team can see those in the document (Figure 8-6).

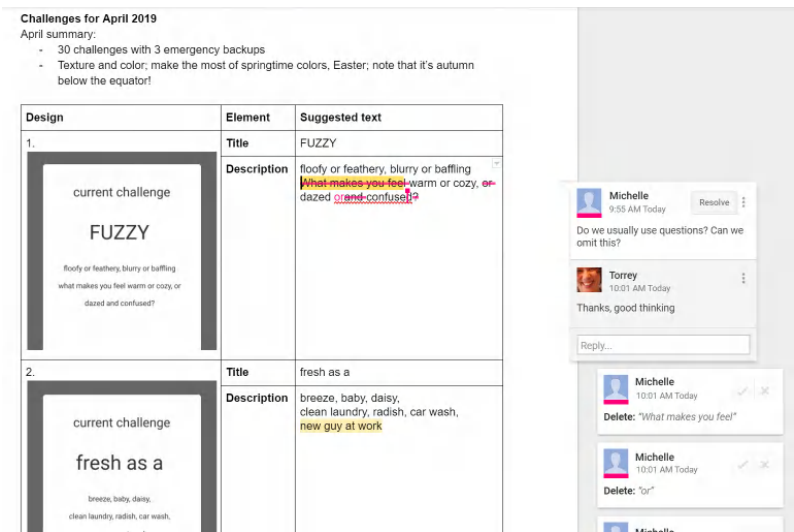


FIGURE 8-6

Michelle has already made comments and suggestions in the review document for April challenges.

When all of the reviews are complete, I work with the designer, developer, or content management system, as appropriate, to publish the final text. The document is kept as part of the online record of the development work. For sensitive topics, like health, privacy, security, and more, the document would be available to be requested by regulators or subpoenaed by courts in the future.

Publish the Text

To become part of the final experience, UX text must become part of the code. Usually, that means a work-tracking item is assigned to an engineer, with the final designs and text review document linked to

that item. As a feature is developed, the engineer may also add error messages as the need becomes apparent, sometimes unbeknownst to the content designer or product owner.

REVIEW TEXT IN CODE

When the code for a feature is ready to review, error messages and all, the content designer may receive a code review request. This is our best chance to check for typos before they go into the build as well as to check that the designed text was entered correctly. If there are new error messages, I'll check them and suggest alternative text if necessary. I'll also check that the engineer's code comments reflect the purpose of the content.

For content designers without development backgrounds, the code might look intimidating! Add to that the complexity that the same experience may be coded in different languages for different platforms, like iOS, Android, and the web. In general, text is coded into a single file. Each “string” of text is there, plus a comment about its purpose, context, and any special accessibility or localization instructions. For example, let's consider the text shown when 'appee is not connected to the internet (Figure 8-7).

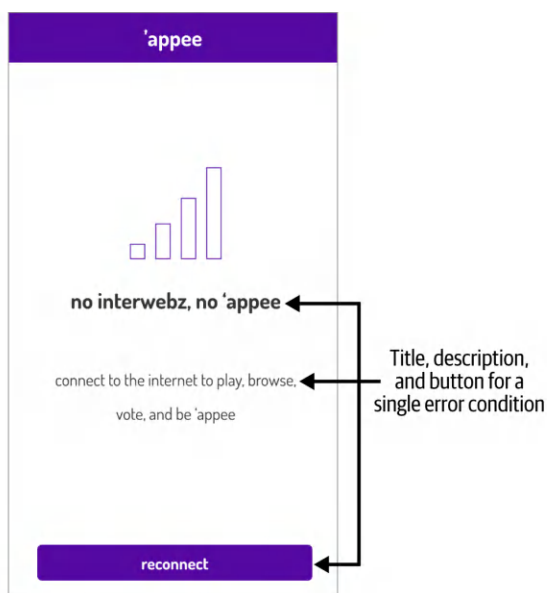


FIGURE 8-7

A single error message within 'appee requires a title, a description, and a button.

The text may be coded in XML for Android and in Swift for iOS. In XML, those lines of text might¹ be coded like this:

```
<string name="errorConnectionTitle">no interwebs, no  
\'appee</string>  
  
<string name="errorConnectionDesc">connect to the  
internet to play, browse, vote and be \'appee</string>  
  
<string name="errorConnectionButton">reconnect  
</string>
```

In Swift, those same lines of text for the user experience would appear differently. They might be coded as this:

```
String(localized: "E_CONNECTION_TITLE", defaultValue:  
"no interwebs, no \'appee", comment: "title for no  
internet connection")  
  
String(localized: "E_CONNECTION_DESC", defaultValue:  
"connect to the internet to play, browse, vote and  
be \'appee" comment: "description for no internet  
connection")  
  
String(localized: "E_CONNECTION_BUTTON", defaultValue:  
"reconnect", comment: "button for no internet  
connection")
```

To review text in code, I need to be able to find the text and review it and either correct it or flag it as a bug for the developer. Note that the same text that was written in the design (Figure 8-7) is here in the code, but has some changes. For example, the apostrophe that begins the app name (‘appee) has a backslash in front of it. This is called “escaping” the character, so that the software will understand it as a character instead of as part of the programming language.

If you’re going to design content for user experiences, it’s useful to know how that content is coded for publishing. Learning the basic grammar of the programming languages for your experience lets you catch problems before they go live, usually with just a few minutes of review for entire features of text. But if you’re very lucky, your organization will invest in a content management system for UI text.

1 Please don’t take this as a recommendation for any particular coding conventions in either programming language. To learn how to code, please see other books in O’Reilly’s extensive catalog!

CONTENT MANAGEMENT SYSTEMS FOR UI TEXT

Alternatively, some organizations have invested in content management systems (CMSs) for UI content. UI-centered content systems, like Ditto and Frontitude, allow the content to be reviewed and reused (and sometimes added and updated!) within an interface built for nondevelopers. I have used content publishing interfaces where the writer tasks are as simple as entering text into input fields or as complex as coding custom XML. These tools can also be integrated with internationalization systems so that translated or localized content can be added to the code without requiring developer intervention.

Having a CMS for UI text can allow content designers to streamline the work they do across the software development life cycle, from initial designs to time-sensitive content delivery, to bug fixes and late-breaking error conditions. Having a single tool that captures the current state of the content and the history of content changes is unusual, though, and it can't capture the work done outside of those pieces of text, from design iterations to system-wide content work that meets strategic goals.

Track the Content Work to Be Done

There are very few people on any team who work throughout the software development life cycle the way a content designer does. There often isn't an obvious way to become aware of all of the work to be done, or to keep track or prioritize it.

When there is so much work to be done and that work is distributed across so much of the organization, it can even be difficult for the leaders in the organization to understand how much work is getting done. The easiest way to show them is to track the work from the beginning.

I prefer to use a work item, bug, or ticket-tracking system, like Azure DevOps or Jira. As long as the tool will hold the data we need, it will work. It's most useful if the engineering, design, support, and UX content team are all on the same system so that we can pass tasks back and forth to work on them. But if the teams are inconsistent in the systems they use, it's OK to set up one specific to UX content.

The most basic information I need in any given work item is the following:

- Definition of the task
- The problems the task solves, including how we'll measure success
- Priority
- Current status
- Files or links to files
- Date the ticket was created
- Date the last change was made

To provide the simplest example, suppose that I work at 'appee and I find out about new UX content: 'appee will be adding a new direct-message feature for people playing in a challenge. I open a work item to track the UX content work for that feature.

When the product owner has a meeting to kick off the work, I start a new document to take notes. I add a link to that document in the work item. I also add links to the product owner's document. At the end of the meeting, I add pictures of the whiteboard to the same work item.

I also open work items to track work that I initiate as a UX content project, like the rewrite of an existing UX flow. Whatever work is needed, it's a matter of seconds to make a ticket for that work, and I know that I'll keep track of it.

I can assign each work item to the team member who needs to take action on it, whether it's the legal partner whose approval is needed or the engineer who needs to code the text. When it's assigned to me, I know that I have work to do; when it's assigned to them, they know it's their turn.

Tracking the work allows a single, reputable answer to the question "How much work is there?" When business decision makers, department heads, and product leads ask where UX content is most needed, I can sort the work items to let me answer with the number and priority of work items. The tracking system lets me sort work items by status, priority, which person it's assigned to, the date it was created, or any other criteria in the list. I can use the links and content attached to the work items to organize the work and remind myself about what's needed next.

Similarly, tracking work items helps me tell the story of the content team’s contributions and speak to the impacts that we make. At the end of a project or review period, the accomplishments are easy to summarize by team and priority, to showcase their impact.

When the team understands how much content work there is (and how valuable it is!), it can be difficult to prioritize all of that content work to be done. Even when the writing work is tracked and the process is established, it can be difficult to make sense of what you should do first, or even next.

Prioritize What Is Urgent and What Is Important

I like to use the *Eisenhower Matrix* for UX content tasks, which categorizes work according to importance and urgency (Table 8-1). Any task or work item is either urgent or not urgent, depending on the time frame you’re considering. At the same time, the item is either important to you or not important to you (but might be important to other people in your organization). Those four categories come with implied actions:

- First do the work that is both urgent and important.
- Schedule important but nonurgent work for later.
- Delegate the urgent but not important work to people who find it important.
- Discard the work that is neither urgent nor important.

TABLE 8-1. The Eisenhower Matrix, as applied to UX content tasks to be done

	URGENT	NOT URGENT
Important	Do Design new experiences Unblock design, engineering, research Write text that affects liability	Schedule Repair of existing, broken text Research into effectiveness and usability Updates to voice and terminology Partnering about design strategy
Not important	Delegate First drafts of common, edge-case, or error text	Discard Arguing about grammar, like prepositions at the end of sentences

Work that is both urgent and important should be prioritized over any other work. This includes work that other people are currently engaged in. Developers may be coding new experiences or updating current experiences, or they may have just uncovered a failure case for which they need a new error message. This also includes future-facing design and research—keeping designers and researchers unblocked. Designers should have the best possible words before their designs are reviewed, and long before coding. Researchers should have the best possible words in their usability and concept studies to evoke the information that will be the most useful later on.

When work is important but not urgent, we can track it in the work-tracking system and schedule time to do it. This work includes all of the content that content designers recognize as broken but nobody else is working on. We can make time to create new content for those experiences and lead those projects. These changes are not to be made lightly; we will need to communicate the changes we want and the impact we expect those changes to make. Part of the work will be to articulate how the content underperforms now and how we will measure the effect of the changes.

When work is urgent but not important to us, that means we don't expect it to significantly help us reach the goals of the people using the experience or the goals of the organization. This content work should be delegated to the team member it is most important to. This might be the first draft of a rush-out-the-door experience, notification, or message.

Encouraging other people to do the initial content design might seem strange, but it's a great way for teammates to express what they need out of the content. They can leverage existing content design system tools, working with voice and tone, terminology, style guides, text patterns embedded in the design system, and LLM-based tools, and reuse content written for elsewhere in the experience. If necessary, we can provide a reduced level of support after they have self-served. In that case, we glean what they intend, clarify with them, and help them simplify it. It can save time for both of us and give us the opportunity to build a more solid partnership.

When work is neither important nor urgent, it's OK to not do it at all. This includes almost every argument I've had over grammar, commas, and hyphenation, except where the text change would change the

meaning of the phrase. Arguments are an important part of the mix of communication in a healthy team, as long as people are basing their arguments on how to best meet the goals of the organization and the people who will use the experience. But even more important is to build systems of content design, including processes, in which the person responsible for the words can be trusted to make this kind of decision.

Systematize Content Design

One of the most important tools in a content design toolkit is the systemization of the text used throughout the experience. Similar to visual design systems, or even backend engineering designs, a content design system has to be built and maintained as infrastructure. It needs to be designed in anticipation of the stresses it will experience: to change, be ignored, or undermined by people who don't understand it.

When the content design system is adopted within an organization that already has content designers, two things happen immediately:

1. The workload of the content designer is reduced for low-impact tasks, because the team starts to self-serve using the content design system. They can then focus on the higher-impact design work to be done, for example, making sure that new features will be highly usable at launch.
2. The technical value of the content designer is better understood, because the team starts to realize how much work goes into choosing the right words.

The longer-term impacts of the content design system can be realized in faster adoption, lower support costs, greater retention, reduced time to design, and greater brand affinity across the product. That's a big promise, I know! But it's a safe bet: these are all fundamental UX promises that require usable, brand-aware UX content. If your organization has been making these things happen without a design system, you've been reinventing the content design for every feature, every time. Creating a content design system will at least save the time and effort required to design, allowing the team to realize all of those other benefits faster.

When an organization is just starting to think about systematizing their content design, it can be daunting to find the right place to begin. I propose thinking about that system in four parts: principles, personality, patterns, and practicalities.

PRINCIPLES

Principles define the direction the design should go, at the highest level. These probably already exist for your organization. They are only rarely appropriate to create in isolation from the rest of the product or service design

I've seen design systems include values, principles, and even mission or vision statements. What's important for this section is to realize that principles aren't prescriptive about a particular solution, but instead talk about the characteristics or impact of the design of the final experience. This part of the design system includes any attempts to articulate the direction to take design work.

Product and design principles, sometimes called *product values*, may be formalized at the company or product level, within a design system, or within organization documents. We discuss product principles in Chapter 2 as the basis for the voice chart. Design principles, in contrast, may include general principles of good design or specific design characteristics that the organization wants to use to achieve its goals.

I also include audience definition work, including personas or jobs to be done, in this overall “Principles” section. Who we're building for matters. It will affect the language we choose to address them with—from the words and grammar we choose to the complexity and metaphor we might allow.

As we start to define the solution, however, we start to move into other parts of the content design system.

PERSONALITY

The voice chart discussed in Chapter 2 is one common element of a content design system. It is built from the product principles and belongs in the personality. It provides a definition of the brand voice, a place to start for content iterations, and a way to tie-break decisions about what text belongs in the product.

If you have a chatbot (whether or not it uses AI) in your product, especially if it appears as an interactive entity, it also requires a personality definition. This is sometimes called the *persona* of the chatbot. It needs to include the entity's backstory: where it came from, what pronouns it uses for itself, how it refers to others, what it values, how it will respond to rude or inappropriate suggestions, and more. This persona might

include an articulation of the chatbot persona's goals, motivations, and activities, or any other information that would help the chatbot's designers continue to create or govern the content it creates.

When an experience interacts with people using a synthetic audio voice, there are many other characteristics to be designed and defined: its timbre and accent, prosody, and many more that are outside the scope of this book. If these characteristics weren't designed, the experience would quickly lose the trust of the people interacting with it, as it would change how it talked depending on which engineering team had updated its data.

PATTERNS

One of the most powerful ways to enable consistent, fast design is to invest in formal patterns. In a design system, this is the systematization of the visual characteristics and interactions possible in the design: the colors, fonts, layout, controls, animations, and more. The graphic and interaction design decisions are formalized and collected into libraries of *components* that make it easier and faster to do design work consistently. In maturing systems, those components are created not only in the designs but also as snippets of code to make it faster to develop working user interfaces.

Design components

LitMop has a library of components, such as titles, labels, and descriptions that can be picked up and dropped into designs, already styled with the correct fonts, colors, etc. It also includes the visual characteristics of buttons, like corner radius, drop shadow, color, and outline. It may include multiple variants of the component, such as different button styling for when a button is active or inactive, and any animation when the button is selected.

The smallest components can be combined into larger components, like the notifications that can appear on the LitMop dashboard (Figure 8-8). The notification component includes the text patterns that have been designed for these notifications. Note that they're slightly different from the generic text patterns provided in Chapter 4: these text patterns have been designed for this component and this product.

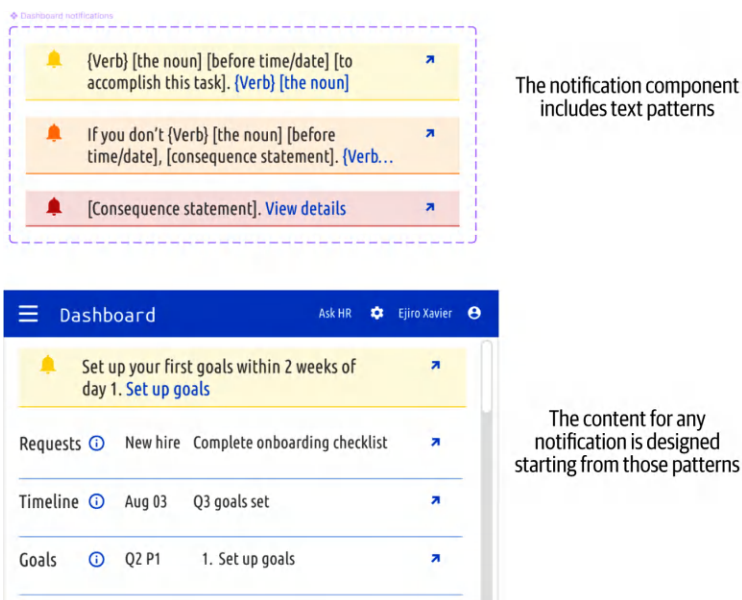


FIGURE 8-8

The dashboard notifications component for LitMop includes text patterns.

Another important pattern within the content is the terminology. A *term* is any word that has a special meaning that differs from normal usage outside of the experience.

Terminology

As I discuss in Chapter 7, a term is a word set aside by the experience as having a specific meaning. Terms need special treatment in UX content and must be systematized so that the same idea is always called the same thing. To ensure this consistency, the team needs to have a list of these terms that includes how they are defined and used. This list is often called a *glossary*.

The casual game 'appee requires only a limited glossary, which can be held in a shared spreadsheet (Figure 8-9).

	A	B	C	D	E
1	Term	Part of speech	Definition	Suggested use	Notes
2	play!	noun	Feature in which people can upload images for a particular challenge	Use the exclamation mark whenever possible.	
3	save	noun	An image the player has tagged to appear in the "saves" area. Does not download the image.		
4	save	verb	The act of tagging an image to appear in the "saves" area.		
5	saves	noun	The area in which images tagged "save" appear for that player.		
6	bookmark			do not use	Use "save" instead

FIGURE 8-9

The terminology spreadsheet for 'appee includes the term, part of speech, definition, suggested use, and additional notes.

Every term exists in the spreadsheet, but not every word. The glossary is most useful when it holds only the terminology that has a special meaning or purpose within the experience, with its part of speech, definition, and any notes about its usage. It's a best practice to include any terms that people are likely to use, even if they shouldn't be used—these can be marked “do not use” with suggestions for alternatives.

In contrast to the relatively simple terminology list for 'appee, LitMop has many more terms that have specific meaning: not only within the experience but within employment law. It's more important for LitMop to have a glossary of terms that everybody can access. There are specialized terminology tools, such as Acrolinx, that can store and provide governance tools.

When terms exist only within the context of a feature, choosing those terms—giving them names—has special significance. It's helpful to systematize that naming process.

Naming

Choosing the names for features, functions, and labels within the experience is a critical design process. When names aren't designed, they can introduce confusion or concern.

Designing names involves research to find out what people already call similar features or functions. This can be their names within competitive products or the names people use to perform the same function

outside of software. For example, people know what “save” means outside the product, but might not know what “tag” means. The best choice for ‘appee is “save.”

Naming processes usually need to go beyond research about what similar functions are called. Stakeholders within the team or organization can have visceral responses to particular names or labels, which may or may not be relevant to the design process. To set expectations and keep the naming process moving forward, I recommend establishing a process for naming before it is needed. Processes that I’ve used at Microsoft and Google include identifying the need for a name, recommended research to perform, a list of criteria for the name, and a workshop structure that invites stakeholders to brainstorm a broad set of possible names. After the broad set of names is reduced to a small set of possible names that meet the criteria well, a trademark search is conducted and the term is localized.

Beyond these patterns that are designed for specific areas and features, there’s also a broader set of practicalities that can be systematized.

PRACTICALITIES

We’ve discussed individual terms that have specific names and definitions, with special ways that people will use them. But the way words are spelled, sentences are punctuated, and numbers are represented can vary widely within a single language. When these words and phrases are inconsistent, the experience seems *messy*. That messiness can create a sense that the product is unprofessional or even untrustworthy.

Some natural languages, like French, are standardized, controlled, or governed. Others, like natural English, are constantly mutating. One way to combat the sense of messiness within an experience is to adopt or create a *style guide*. In English, I’ve used the *AP Stylebook*, *MLA Handbook*, and the *Microsoft Writing Style Guide*. Note that all of these are imperfect for UX writing, but they are extensive! They specify spellings, such as preferring “cancelled” to “canceled,” even though both have the same meaning.

Other practicalities might include the abbreviations to use for the names of countries, regions, units of measure, and other terms within the experience. They can also specify the preferred units to use, for example, the International System of Units (SI) versus United States

customary units. If the experience reports dates, times, or intervals, then specifying how to show dates and times can reduce cognitive load by representing them consistently.

Some tools are available to apply this practical guidance more easily, including tools with LLM features, such as Writer.ai or Grammarly. They can be programmed to detect deviations from the designed style guide, and suggest or automatically correct the text to use the preferred spellings, grammar, or terms. These tools can be very handy! They can also be incorrect and usually require oversight.

Summary: The Tools Are a Means to an End

There are a variety of tools for content designers to use to create excellent UX content. But mastery of any of these tools isn't the point of design. The use of Figma, LLMs, Photoshop, Canva, Google Docs, Jira, or Excel, even at an expert level, won't make you a better content designer.

Instead, we must be willing to use the tools we have to bring empathy and analysis to each interaction a person will have with the words in an experience. We need to draft, edit, and iterate the text using our skill with language to unlock the potential of the experience. We must shepherd stakeholders through a review process to make sure we meet the goals of the organization and of the person.

A 30/60/90-Day Plan

Nobody plans to fail. They just fail to plan.

—UNKNOWN, ATTRIBUTED TO MANY

IN THIS CHAPTER, I DISTILL and explain the plan I used as the first content design employee in multiple teams, of multiple sizes (from 20 to 400 people), at three different companies (Microsoft, OfferUp, and Google). Each time, I joined the team with some idea of the opportunity and a glimmer of the problems I would face, gleaned from the conversations in my interview. Each time, the team brought me in because they realized they 1) had a problem with words and 2) knew that they couldn't fix it without extra help.

The actual number of days is an estimate, not a rule, but they have been pretty accurate for me. Most usefully, the 30/60/90 structure creates three phases in which ramp-up work can be done thoroughly but quickly—and definitely not perfectly. This method helps me to create a basis for collaboration and iteration from which I can not only do the content design work that benefits people using the experience, but also make broad impacts for the business.

Days 0–30: What and Who

The first 30 days are all about learning the experience, the people who will use it, and the team of people who will build it. To be successful, you need to know what's important to each of them. At the same time, you need to build the team's confidence that the time, energy, and money they spend on content will pay off.

Your first task is to find a few teammates who will give you the widest possible perspective on the organization. Those two or three teammates have a couple of key characteristics: they have broad knowledge of the organization, and they know why the organization decided to “fix the words” by hiring you. In the best-case scenario, they also have different points of view from one another.

You should ask these key teammates the following in one-on-one, face-to-face meetings: who is on the team? That is, who do you need to work with to affect what people will encounter in the experience you’re making? Write down names and roles from marketing, design, and engineering teams, as well as product owners, program managers, support agents, forum moderators, trainers, attorneys, business analysts, and executives. In the meetings, try to draw a diagram of the organization and ask these key contacts to correct your diagram.

RELATIONSHIP BUILDING

You should request half-hour meetings with each of those 5 to 20 people you’ve just found out about. (This is very different from a consulting relationship, where access and time to be spent on this kind of discovery are limited. As a new employee, you need to build a different relationship.) There are two purposes for these meetings. The first is to gather information about the organization, product, goals, and customers. The second, and just as important, is to introduce people to the idea of working with you. You also need to understand these partners’ needs, personalities, biases, and drivers. This will help you predict where they’ll help you accelerate the work, and where they’ll stop or delay progress.

Even as I write this second edition, most of the people who make software have never worked with a content designer, and even fewer have worked with a content designer who understands that the purpose of their role is to meet business goals while serving the people who use the experience. By meeting one-on-one, you get the opportunity to introduce them to what it will be like to work with you, and with somebody in our discipline.

In your invitation, you can write something like, “Hi, I’m the new content person on product X. Your name came up as a person who’s important to the product and team, and I’m hoping to learn more from you.”

Next, choose a time that’s likely to be convenient for them, making sure that you have enough time in your calendar to consolidate learnings between meetings.

To prepare for the meetings, make a mostly empty document, whether that’s a slide deck or text document. You can use headings to create structure and then fill in the information that you know so far (Figure 9-1). I recommend keeping the document rough and unpretty to make it clear that you’re spending time learning the information, not “polishing” the presentation. It should also be shareable, and you’ll want to make the link or document available to the people you meet with.

Before the meetings, try to add the information that you think you know, in the briefest, most scannable form. Where you don’t know anything yet, leave the slide or section empty. By doing so, you communicate 1) what you want to know, 2) that you know that you don’t know it yet, and 3) that sharing these things will be valuable to you. Then, you’re prepared to not only take notes in the meetings, but also organize and give context to the information that you’re getting.

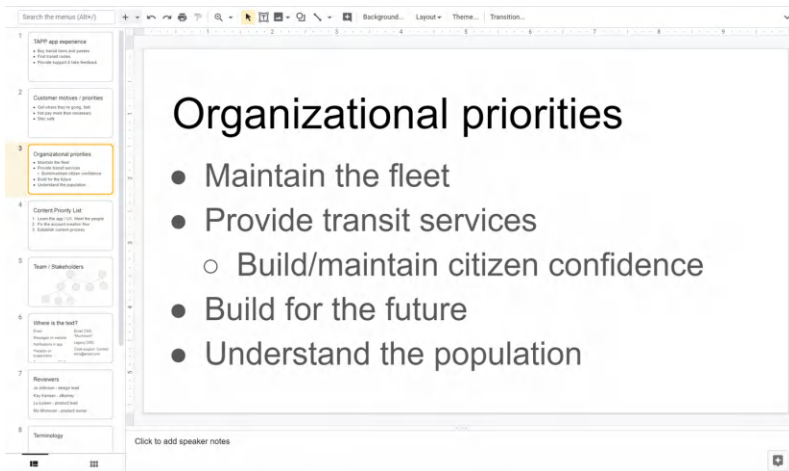


FIGURE 9-1

These notes are structured, as an example, in a slide deck format.

For example, my initial notes have separate sections about the experience, the customer and organizational priorities; the initial content priorities or tasks; a diagram of the team and stakeholders; and places to add info about the channels, platforms, terminology, and reviewers. I recommend that notes contain at least the following sections:

1. Definition of the experience
2. Customer motives
3. Organizational priorities
4. Priorities for content strategy
5. Team/stakeholders
6. List of existing content
7. Reviewers
8. Terminology
9. Resources: links to existing information

During the meeting, the important thing is to start to build a working relationship with the person you're meeting with. You want to show them that you're interested in the experience, customer, business, and priorities. If the other topics come up, listen, take notes, and move on.

Here are some example questions that I ask:

- What's the most important part of the experience?
- Who are the customers? Is a customer someone who installs, someone who uses, or someone who buys within the experience? If it's an experience people use at work, are the customers the people who buy it or the people who use it?
- How do these people solve the problem right now? How is that experience different?
- What's important to them? What motivates them? What are their priorities, their desires? Do we know what they like or dislike?
- Among the people making and supporting the experience, who will be an ally in making it a great experience? What are their motivations, hopes, desires for it?

- In the organization or industry, is there anything working against us? Is there anything working in our favor?
- What's the most important thing I can work on?
- Where are the words broken, or where can the words help the most?

As you listen and learn, you present the document and take notes at the same time, as much as possible. That way, you can show in real time that you're adding that person's priorities to your priority list, and adding their information to your understanding. If what they say is already represented, ask them to check and correct what's there.

Between meetings, consolidate what you've learned. Note-taking can become very messy! Sometimes I add notes directly to my document, and sometimes in comments. Other times, we use a whiteboard or paper in the meeting, so I take pictures of the notes to consolidate later. Often, there are existing links or documents the person wants to share, which can be saved in their own list. When we're in an online meeting, I can record it to refer to later.¹

Some of the most valuable knowledge that comes out of these meetings is the list of existing content. In my experience, if a team has been working without a content professional, nobody actually knows what all of the content is. No single person has a coherent view of all the content that a person using the experience might encounter.

So, when somebody mentions a folder, repository, content management system, or other source of UX content (for example, UX text, help content, social media engagement, emails, notifications, websites, or canned responses), I add it to my notes. Any content that affects the person's experience is part of the content story. It's content I should be aware of, even if I never work on it.

Similarly, my ears perk up whenever I hear words used with special or unusual meanings. I add those to my notes as a nascent terminology list. As I build the list of terms, I attempt definitions of those terms. When those terms come up, I ask teammates to check and correct my understanding. By using a common, sharable tool to create clarity for myself, I also help the team create a common understanding of the terminology we use.

¹ There are a lot of virtual note-takers that can provide summaries of online meetings, which some people find very helpful. I do a lot of my thinking and analysis as I formulate my own notes, so I prefer not to use those tools.

EXPERIENCE LIFE CYCLE

As your understanding of the experience begins to mature, try to draw the life cycle of the experience. You can begin with the cycle in Chapter 1, Figure 1-9, and adjust it for the experience you’re working on with this team. You should show the journey of a person through the cycle of investigating, verifying, and committing to trying the experience, then setting it up, using it, and hopefully coming to prefer it. Then, you can adjust the length of sections to reflect the reality of this experience, this organization, and the people who will buy and/or use this experience, at this moment in time.

For example, Figure 9-2 shows a diagram of the virtuous cycle adjusted to show the TAPP experience. TAPP attracts people when they are interested in using transit, converts them by providing routes and fares that will work for them, and onboards them by selling them fares and passes. They start to use TAPP with their first ride and might engage to the level of using transit regularly. A person could become so enthusiastic about their experience with TAPP that they become an advocate for using transit and bring other people into using TAPP.

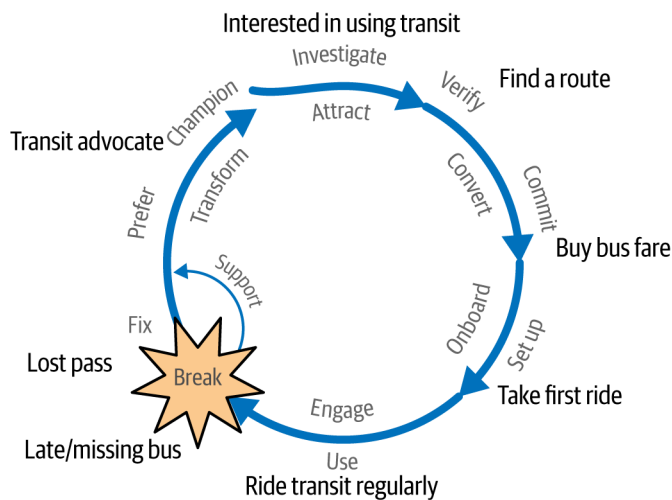


FIGURE 9-2 The virtuous cycle diagram has been adjusted to show the TAPP experience.

When you have the experience drawn, add it to your notes. As your meetings continue, you can use the diagram to ask members of the team where they think the experience isn’t working. You can also use it

to explain what you're there to do: you will be making the content that will help spin the wheel for the organization and for the people who will use the experiences you make.

START WRITING

By the end of the second week, if not earlier, you might begin to receive tactical requests to fix the words: "Can you rewrite this email? What should go in this error message?" I start these first writing tasks in parallel to my learning about the organization and the experience, because the strategy will work only if the writing can appear in the experience.

These first writing tasks are a great testing ground for the ideas percolating in your brain about who the people are that will use this experience, what your organizational purpose is, and how the person's and organization's priorities could be expressed in the UX.

This is also an opportunity to demonstrate how you work: asking questions about goals and purposes, measuring success, and drafting UX text in the designs. Possibly, you just have a screenshot of bad text that you will edit to provide different text (for example, see Chapter 8, Figure 8-1). The person requesting the text might expect nothing more than an email or chat message with the new words to use, but you have an opportunity to demonstrate that UX text should always be reviewed as part of the design, the way the person will encounter it.

Part of designing is finding out how your new team uses design tools. This might be your designer's first time having somebody else working inside their files. It's a good time to talk about how you need to work: in the designs. You could share your own preferences, for example, whether you prefer to do initial drafting in a copy of the design file or work on iterations within a shared file. Your group may already have standards for design file organization and how a design system is used, and have existing resources about terminology, voice, tone, style guides, brand guidelines, or principles, or they might not exist at all.

Once you get started, I encourage you to write at least three good options for the first content you deliver. Make each version meet the purpose the person will have for that UX text and the purpose the organization has for that screen. Strive to make the options as different from one another as possible. By presenting a few, curated options, you

enable the team to have a good conversation about the purpose for the UX text while you build their understanding about the power of what you can do with words.

To the person who requested the text, you need to explain the reasons why any one of the three options might be the right choice. Often, I learn more about the problem at this time, and I need to draft more options! This revising is a normal part of the process, and it lets me understand the experience and the organization at a practical, hands-on level.

When you and the requester agree on one or more of the text options, ask the following: “Who else should be reviewing this?” You might suggest some names that you learned from your interviews and use the names they recommend. Send or share your first requests for review, listing your recommendation first, and one or two alternates, including the reasoning.

At the end of these first 30 days, you will have talked to most of the right people; you’ll be in most of the right meetings and internal communication channels, like group emails and chat groups; and you’ll have drafted your first text.

At this stage, the document that you started at the beginning of the month now contains at least the following:

1. A prioritized list of tasks to produce or improve UX content
2. The motivations and priorities of the people who use the experience
3. The organization’s priorities and constraints
4. Beginnings of lists: channels, terminology, content reviewers
5. Links or images of initial, tactical content work

You’re ready for the next phase after you have built these new relationships with your team and equipped yourself by gathering the information in your notes.

Days 31–60: Fires and Foundations

In this second phase of work, half of your time is spent chipping away at urgent, “on fire” work. Doing the “on fire” work helps provide a basis for you to test, practice, and create the foundational pieces that will help the work go better and faster in the future. It will help you to build your

understanding of the team, the experience, and the people who will use the experience. Just as important, it helps build trust with the team when they see you assessing the problems alongside them, as part of the team.

As much as possible during this second month, you should delay effort on larger, systemic changes. The UX text that you write in this second month is unlikely to be the best writing you'll do for the experience. It won't be consistent, because there's no consistency defined. It won't be in the ideal voice, because voice isn't defined. To do good work on systemic changes, the work must be strategically aligned and prioritized with the development schedule to avoid randomizing the team and fracturing your own attention.

Right now, before those systemic changes begin, is the time to measure the baseline of how the UX content is meeting the goals of the organization and the people who use the experience. It's time to examine the "broken walls" in the experience. If the team can't tell where in the experience people drop out, or where they fail to engage, or where they make the decision to buy or commit, now is the moment to specify and advocate for the measurements, research, or instrumentation necessary to notice a change.

You also should try to use the experience yourself, recording the experience and taking screenshots as you go. If possible, you should consume the usability research already conducted, if it exists. Then, apply heuristic measures to key UX flows in the existing content.

You now need to make an initial report on what you find, including what you know about the behavior and sentiment of people using the experience, and a scorecard of content usability based on the heuristics. These initial reports indicate what's working, what isn't working well yet, and which work you recommend prioritizing.

At first, you can share the report informally with the members of the team most directly involved with creating the experience. This report outlines problems in what they have already built, so you don't want to share it too widely or with too much fanfare. Everybody who built the experience before you arrived was doing their best and working hard, and you need to respect that. Later, after the experience is improved, you can use the report as a baseline from which to measure improvement.

While I work on individual content requests and measurement (the “fires”), I also spend time setting up the foundational pieces so that I can work and collaborate faster and more effectively. I set up tools for content creation, sharing, and organization; code environments; partnerships and processes that integrate with the team; and to track, manage, and prioritize the work to be done.

TRACK THE PIECES AND THE WHOLE

You ended the first 30 days with a basic sketch of content work to be done, and work requests began to flood in. Some requests are for single pieces of UX text, and other requests encompass vastly more: text through an entire experience, error messages, articles and videos, notifications, and more.

On any given day, you might create and review content with designers, researchers, executives, support agents, attorneys, and in code, on multiple projects. You need to learn how this team uses a tracking system or set up a tracking system (see Chapter 8) to serve as a central place to gather, prioritize, and organize information for UX content tasks. By using a tracking system, you keep yourself and your team (when there is a team) afloat on the flood of work to be done.

When all of the UX content work is tracked, you can understand the scope and shape of the work at a glance. You can see where most of the work is concentrated, and find out which parts of the organization you haven’t engaged enough with and which parts of the experience you haven’t examined. Finding most of these blind spots in the second month can help you to avoid problems in the future.

MINIMUM VIABLE PROCESS

Knowing the work to be done gives you the scope of the battle, but it doesn’t help you fight it. In these second 30 days, you need to have tool-chain conversations with engineering, design, and product teams to learn the process for UX text publishing and code review. The first few projects from the first 30 days can help set the context: did that delivery method work for them? Do they have any feedback? What’s the best way to get that person’s input?

You'll need to listen to what people expect and need from you, and what tools they expect you to use together. You need to steer the work toward a repeatable process so that not only is your work easier, but all of your stakeholders know what they can expect from you as well. You should advocate for collaboration tools, seeking the simplest possible process.

Solicit feedback from product owners, marketing, and business leaders about the process. You'll want to ask where it should change to best fit with their system. You should help them understand when and how to involve you, and when and how you'll involve them. Try to find out how you should assign work items to engineers and how to use their code review system.

In every team I've worked in, a decision maker in the organization will respond to the new UX content focus by saying, "I want to review every piece of text." In my experience, they have thoroughly meant it, and they also don't want to sit down and walk through the code every time there is a change.

What they have wanted is confidence that the text won't increase the organization's liability. They want to be sure that the text accurately reflects the organization's brand. They want to have the gut feeling that the words "feel right" in the experience. You want to involve this person in the UX content process right away so that they can resume doing their own job without also trying to do yours. It's helpful to draw a basic content process for them (Figure 9-3) and suggest the moments in that process where you will proactively consult them. It gives you something to point to and say, "Right here, let's look at it together. You give me feedback, and I'll make it right."

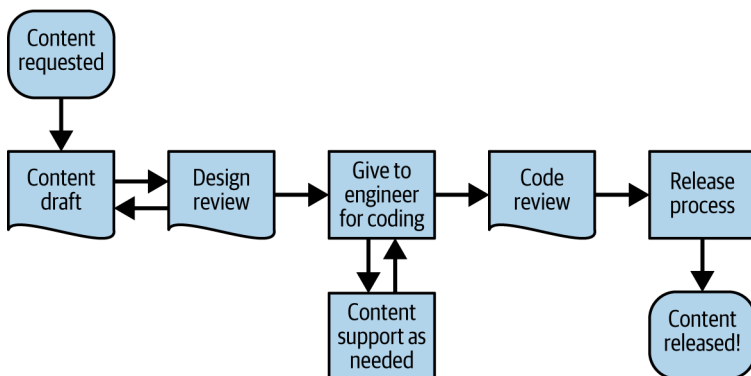


FIGURE 9-3

This diagram shows a basic process for UX content from request, through drafting, review, coding, code review, and release.

DOCUMENT THE CONTENT STRATEGY

To do great work, you need to think systematically about the deep connective tissue of the content: the core terminology and the voice that permeates the conversations the experience has with the people who use it.

Not only do you need these fundamentals of content design, but you also need your teams to understand the systemic importance of content. Documenting the content strategy helps you to demonstrate how it is useful to the organization. Creating that documentation takes its own time and energy, but it pays dividends later.

The purpose of documenting any internal strategy is twofold: first, to ensure that the strategy is well thought out; writing is a marvelous tool to require that thinking. Second, to make future tactical decisions more easily, faster, and more consistently. I use the content strategy to facilitate decision making in the following areas:

- Documented priorities and goals focus the content by defining the core UX problems that it helps to solve.
- Voice charts guide the direction of content creation and iterations and serve as a tie-breaker between good text options.
- Terminology lists create consistency and reduce time spent rehashing choices about how a particular concept is represented in the experience.
- A ratified list of UX content reviewers includes and excludes the appropriate people in the content process strategically, instead of creating tactical political complications.

These are living documents, growing out of the notes that you began in the first 30 days, and they need to be further developed. You should review them on a regular cadence (at least annually) and update them when there are organizational changes.

How will you know when you're at the end of the second phase?

You should have passed at least these milestones:

- New content created
- Tracking system and process established
- Poorly performing content updated

- Legal sign-off on some liability-sensitive content
- Marketing sign-off on some brand-sensitive text

Key indicators of trust should also be evident:

- Leaders are including me as the responsible party for UX text
- Casual requests to work on individual pieces of text
- Active inclusion in early design thinking

About 75% of this strategic work is complete, in a first iteration:

- A tracking system for content tasks
- Alignment about motivations and priorities
- Knowledge of the existing content and how to access it
- Terminology list
- Voice chart

Another indicator for me that I'm at the end of this phase is a feeling that I'm finally set up to accomplish more. Now, the most urgent, distracting, tactical work is complete, and the important foundations are laid. The UX content is ready to have an enormous impact on the quality and effectiveness of the experience. It is time to move forward!

Days 61–90: Built for Growth

The strategy is as done as it can be for now, which is as good as a strategy ever gets. Now it is time to build a presentation about the strategy as a whole for the first time and share that new foundation. The desired outcome of the presentation is to cement the solid foundation: your team and your leaders can have confidence that the content strategy is created considerately and together, and that the work has a purpose. By signing off on the strategy, they validate and support the work to be done.

Communication about the work is a critical part of the work, and might be the most difficult part to accomplish. The presentation includes all of the parts created to date: the problems it solves, alignment on motivations and priorities, the tracking process and list of current tasks, the content landscape, the terminology list, and the voice chart. This

summary of the UX content strategy is solid enough that the important (or controversial) ideas are covered, but unpolished enough to indicate that time hasn't been wasted fussing with internal documents.

Ideally, everybody at the presentation has participated in the process of creating the strategy. They get to enjoy the fruition of their own work and advice and, accordingly, the results of their decision to hire their first UX content pro!

At this point, you should seek feedback. Feedback provides the corrections necessary now to be successful later. If the feedback is that the strategy is wrong, you should thank them for their perspective. If they're wrong, it might be simply that the presentation of the work didn't scratch their itch. If they're right, it's fantastic that you've gotten correction so early. It's only the second or third month on the job, so it's the best time to make adjustments. Even the presence of feedback is a good sign, because it means people are invested.

Starting in this third month, you need to set the sustainable pace of how you'll engage in the tracked, prioritized work of content design. You should respond to requests and make requests for content changes. You should partner with the team to design new experiences, advocate for the people who use the experience every time they consider the words, and apply and sometimes revisit and tweak the strategy.

This phase is over when the process of creating content for the experience is healthy enough that you can begin to broaden the scope of what the content strategy can do for the organization. Now is the time to check in on trends in the field and on the rest of the content being created about the experience. You'll want to work to strengthen connections with marketing, operations, and knowledge management. You should investigate opportunities in the industry, like content-bots using machine learning to prewrite content, and seek out new research, like best practices about titles, labels, accessibility, and inclusion. Some words might still need "fixing," but from now on, new UX content will be created strategically, the first time.

Summary: To Fix the Words, Build Strong Foundations

To introduce a new discipline into a team, there's a lot of work to be done. I start my 30/60/90-day plan by understanding those goals and my teammates in this adventure, what work they have done, and what work they don't know to do yet. In the second month, I fix urgent problems while building the foundations that will help me organize and demonstrate the effect of future work. Finally, in the third phase of this process, I begin the work that uses the power of content tools to be more effective than ever before.

After these 90 days, I am confident I can more rapidly and consistently create UX content than the organization could before I joined. I can probably demonstrate, also, that the content is likely to be effective in meeting the organization's goals and the goals of the people who will use it.

There are a lot of content accomplishments in this plan, but the work that pays the most dividends comes from doing the work visibly, in partnership with the team. By making the content strategy visible in presentations about voice and terminology, the team and executives understand that we've unlocked a new power tool. By making the content tasks visible, the team understands the work it takes for the UX content to advance the goals of engineering, design, and the organization while supporting the person who will do that work. In the process of doing the work, the UX content goes from being a source of pain to being valued.

Advocate for Content Design

It is unforgivable to do what one doesn't love, especially if one succeeds.

— CHRISTIAN DIOR, FASHION DESIGNER

IF YOU TAKE NOTHING ELSE AWAY from reading this book, I hope you know that the purpose of UX content is to meet two categories of goals: those of the organization and those of the people who use the experience. To meet those goals, you'll need to listen to people, prioritize the work, and collaborate with your team. You'll need to advocate for the concept that content work needs to be done, whether or not your role is dedicated to content design.

Ground the Content in Empathy

When we create experiences, we need to care about the people who will use the experience. When we don't care, we risk failing at our core task: to make experiences that meet their goals.

The root of caring is to believe people when they talk about the experiences they have. Their experiences can be similar to your own, or they might be literally unimaginable, but we don't need to imagine them. We must listen to what actual people say, observe how they behave, and believe that we are hearing their story.

When most people listen to a person's story, they tend to produce the chemical of caring: oxytocin. When content designers listen to a person's story, we get that oxytocin and more.

For a content designer, the simple act of listening uncovers a gold mine. When people tell their stories, they are likely to use the words they will find recognizable. By listening, we learn the grammar that the people already understand. We learn the emotional lading of the jargon specific to the people's experience.

When the content designer then uses these words, they can create an experience that connects people to the experience without feeling like they are reading.

To write effective UX, work toward understanding the concerns, needs, and words of the people who will use the experience. Go out and listen to them. Bring them in and listen to them. Watch videos of interviews with them, and seek to understand their point of view. This research will give us an appreciation not only of where they're coming from, but also of how different our own perspective is.

And while we talk to people outside the organization, don't forget the people on the team. These people, with their opinions, viewpoints, perspectives, and prior knowledge, will have an enormous impact on the experience, too. There are people invested in making a great experience all over the organization, including the marketing directors, general managers, directors of design, heads of engineering, engineers coding the feature, program managers, product owners, designers, and the sales and support agents.

Anybody, and everybody, can have opinions about words. How to use those words systematically to meet specific goals might not be well understood, especially when there hasn't been a dedicated content designer.

Advocate for UX Content to the Team

If you accept a job as the first content person in an organization, they might think that you're there to "choose the right word" or to "check the words." They probably think of it as a word problem: "We need to explain," or "We need them to understand..." Or, maybe it's a UX problem: "We need words to go on the buttons" or "There are too many words on the screen."

"We need words" is not the problem that we solve as content designers. We communicate. We invite action. We inspire loyalty. Our teams need to know that content design can be used to solve problems. It's up to us to frame our work to reflect the problems we're helping to solve.

In 2023, I had the opportunity to gather together with a group of content design leaders. Attendees were asked to bring something useful to share: an idea, a process, whatever. The door was wide open. I offered something audacious: I'd like to lead the collaborative writing of a manifesto about content design.

The job of a manifesto is to compile a set of ideas, defining its own spot in peoples' common ground for discussion, disagreement, belief, or rejection. To get a content design manifesto right, it has to be a collective effort; the practice of content design has evolved rapidly and differently in different contexts. The manifesto is available online and reprinted here for ease of use.

THE CONTENT DESIGN MANIFESTO

Why this document

We want to help everyone understand who content designers are, how we work, the impacts of our work, and the future we plan for ourselves. We hope this understanding results in more success for content designers, for the organizations that employ us, and for the people our products serve.

Who we are

We are content designers, shared owners of user experience (UX) design processes. Our job is to design experiences that are clear and meaningful for the people who use them. Clarity and meaning drive business impact for our organizations.

Our skills and perspectives are unique in our teams, and uniquely enable us to ensure UX content is accurate, actionable, and inclusive for every person—wherever, however, and whenever they need it.

How we work

We work as one of the many disciplines of experience design. Our work is informed by data, research, strategy, business objectives, and deep understandings of human behavior.

We work with comprehensive breadth, depth, and scope across product areas. We connect with customers and collaborate with partners and stakeholders in engineering, research, product management, design, legal, marketing, customer service, and executive roles, among others.

What we know

1. **Content is valuable.** Organizations can't create experiences for people without UX content. Content designers reduce expenses and increase income, impacting metrics for engagement, conversion, adoption, retention, satisfaction, brand affinity, and more. Content designers reduce risk, ensure quality, increase accessibility and inclusion, enable discovery, and affect how people feel.
2. **Content design is more than writing words.** We design across systems to ensure cohesive, ethical, high-quality content that reflects how people think. We create outputs such as brand voice, content design systems, conceptual models, governance processes, content-centered research, and terminology.
3. **Good data enables good content design.** Content design requires a deep understanding of people, their needs, and the business opportunities. Like all other design, content design improves with data-informed iteration.
4. **Details matter.** We apply rigorous thinking to every iota of the content and its communication, because it matters to the audience's experience of the product.
5. **Tools that ease language empower content designers.** We embrace large language models, spelling- and grammar-checkers, and other tools. These tools help us do the heavy, repetitive language work, enabling us to focus on strategic impact.
6. **Words that aren't working indicate a UX deficiency.** Content designers ask important questions, conduct research, and iterate designs to uncover and correct underlying problems in the design and product strategy, even when words can't fix the problem.
7. **Content design works from product concept to product launch.** Our value is in the depth and quality of our thinking. Expect the most value when we are engaged early in the design process and allowed time to work.
8. **Our community is inclusive.** We hire, mentor, sponsor, and educate content designers so that our community reflects and serves the world's populations.

Where we're going

With our truths in mind, here's where we intend to take this field in the next few years. Any one content designer may have few or none of these in their direct control, but as a group, we can influence the industry to shift in this direction by choosing our projects, employers, and priorities.

1. **Content designers, in-house and freelance, will be compensated competitively.** Our skill sets and scope of ownership are comparable to product designers and UX researchers, with different outputs. Many companies have already begun to pay all UX roles equally.
2. **Content designers will drive product design projects.** Leading strategic design is a skill we share with our peer product designers and UX researchers, aided by our skills with language. Some projects are best approached visual-first, and others content-first; the most suitable practitioner will take the lead.
3. **Content design managers will be elevated as cross-functional leaders.** Leading and managing teams are skills we share with design managers and UX research managers, enhanced by our skills with language.
4. **Content designers will selectively focus on strategic, high-impact efforts.** Beyond its cosmetic importance, content design is understood to be foundational to UX. Content designers work end to end as valued strategic UX professionals.
5. **Access to content design careers will expand.** We promote free opportunities for those looking to get their start, make entry-level positions available, attract career-changers, and work to amplify credible resources and certificate programs.
6. **We will mitigate the carbon footprint of digital content.** We strive to balance the needs of our audience with our impact on the planet. Our work contributes to efforts that increase sustainability and reduce harm.

At the time we collaborated, I didn't know whether we would decide to make this document public. But after we drafted it, we recognized the power it could have. Similar to the *Agile Manifesto*, the *Content Design Manifesto* exists to enable people to describe our purpose and power to

their teams and organizations. Within its first year, the *Content Design Manifesto* was signed by more than 1,180 people around the world and made available on the manifest's website (<https://thisiscontent.design>) in 28 languages, all by volunteers. I hope it's helpful to you.

Summary: Use UX Content to Meet Your Goals

Organizations that make experiences are learning the positive effect that UX content can have when it is written strategically. Content designers, people dedicated to creating UX content, can bring knowledge of best practices, UX text patterns, structures for voice, iterative editing, and review.

Perhaps you design content, you support a content designer, or you're considering adding a content designer to your team. I'm so excited about our future. We have solid work to build upon, and we have so many possibilities ahead as we continue to invent and research best practices. Together, we have the opportunity to help people and organizations meet their goals by creating, iterating, and measuring the UX content.

[Index]

Symbols

- 30/60/90-day plan, 233–247
 - days 0-30, 233–240
 - days 31-60, 240–245
 - days 61-90, 245–246
- * (asterisks) in descriptions, 71–72

A

- abandonment, 174
- abbreviations, 231
- A/B testing
 - cost reduction, 175–176
 - defined, 170
- accessibility heuristics, 187–191
- accuracy of content, 128–129
- Acrolinx, 230
- activity, defining, 172
- addiction, danger of, 173–174
- AI (artificial intelligence). *See* *also* LLMs
 - embedded, 114–115
 - writing text with, 215–216
- alternate text, 190
- 'appee
 - concepts, 27
 - confirmation messages, 98–99
 - content review, 218
 - defined, 20–21
 - descriptions, 72–73
 - empty states, 78
 - engagement, 172
 - error messages, 108, 212
 - grammar, 33
 - labels, 83–85
 - measuring behaviors, 169–170
 - menus, 66
 - notifications, 102–103
 - onboarding pace, 171

- product principles, 25
- punctuation and capitalization, 35
- purpose of, 25
- review document, 218
- scorecard, 185–206
- sign-in screen, 40
- text input fields, 90–91
- titles, 60–61
- transitional text, 94–95
- verbosity, 31
- vocabulary, 29
- voice chart, 38
- AP Stylebook, 231
- artificial intelligence. *See* AI; LLMs
- assessments for understanding, 181
- asterisks (*) in descriptions, 71–72
- audio voice, 228
- auditability of content, 121, 130
- autocratic decisions, 44–45
- automatic corrections and suggestions, 215
- autonomous decisions, 43
- Azure DevOps, 222

B

- badges, 8
- Bethbot, 18
- bias in LLMs, 128–129
- blocking errors, 109
- broken experiences, 1, 3, 8, 105
- Built for Growth phase of 30/60/90-day plan, 245–246
- Buolamwini, Joy, 113
- Burnett, Margaret, 198
- buttons, 64–69
 - brevity of, 64
 - matching wording in titles, 67

- order of, 65
 - patterns, 69
 - purpose of, 64, 69
 - tips for, 69
- C**
- capitalization, 34–36
 - card-sort exercise, 179
 - ChatGPT, 51–52
 - clarity
 - actions with unambiguous results, 197–198
 - editing for, 163–165
 - heuristics for, 196–200
 - CMSs (content management systems), 222
 - code review requests, 220
 - codesiging, 179–180
 - collaboration
 - content review, 216
 - design, 48–53, 210
 - comments, 217
 - communities, 10
 - completion metric, 174
 - concepts, in voice charts, 26–28
 - concision
 - editing for, 160–162
 - heuristics for, 193–195
 - confirmation messages, 98–101
 - patterns, 98, 101
 - purpose of, 98, 101
 - tips for, 101
 - consensus decisions, 43
 - constraints, identifying, 11–12
 - content designers
 - content design manifesto, 252–254
 - content design systems and, 226
 - dealing with existing content, 237
 - defined, 251
 - future of, 253
 - listening to experiences, 249–250
 - training with voice charts, 41
 - work methodology, 251
 - content-first design, 47–56
 - collaborative design, 48–53
 - transforming conversation into experience, 53–55
 - content management systems (CMSs), 222
 - content reviews, 216–219
 - content strategy
 - documenting, 244–245
 - presenting, 245
 - context-setting titles, 59–62
 - context, writing for, 209–216
 - AI text writing, 215–216
 - design tools, 210–211
 - lo-fi tools, 211–214
 - controls, 85–88
 - grouping, 86
 - name and state, 86
 - patterns, 87
 - purpose of, 85, 87
 - tips for, 87
 - conversation and conversationality, 47–56
 - collaborative design, 48–53
 - defined, 48
 - editing text to be conversational, 162–163
 - heuristics for, 194–196
 - transforming into experience, 53–55
 - cost reduction metric, 175–176
- D**
- DAU (daily active users), 172
 - decision making
 - autocratic, 44
 - autonomous, 43
 - consensus, 43
 - decision makers, 243
 - hierarchical, 44–45
 - voice charts and, 43–44
 - delegating work, 225
 - descriptions, 69–76
 - avoiding asterisks, 71–72
 - defined, 69
 - groups that read, 70–71
 - keyword placement, 71
 - labels versus, 81
 - length of, 70–71
 - patterns, 76
 - purpose of, 69, 76
 - tips for, 76
 - word order, 73
 - design patterns, 228–231. *See also* text patterns
 - components of, 228–229
 - defined, 57
 - naming processes, 230–231
 - terminology, 229–230

- design principles, 227
- design tools, 210–211, 239
- detour errors, 108
- Dillon, Doug, 57
- Dior, Christian, 249
- direct measurements, 168
- Ditto CMS, 222
- documenting regulated content, 217–219
- Dunn, Beth, 18
- dynamic content, 7, 65, 81. *See also* LLMs

E

- Eames, Charles, 47
- editing, 155–165
 - for clarity, 163–165
 - for concision, 160–162
 - for conversationality, 162–163
 - defined, 155
 - goals of, 156–157
 - process curve, 156
 - for purposefulness, 157–159
 - screenshots, 212–213
- effectiveness of content, 19, 167–208
 - direct behavior measurement, 169–176
 - heuristics, 182–208
 - methods for researching, 176–181
- Eisenhower Matrix, 224
- embedded AI, 114–115
- empathy
 - error messages, 106
 - grounding content in, 249–250
- empty states, 76–81
 - in forms, 79–80
 - in menus, 77
 - opportunities for education, 78–79
 - patterns, 80
 - purpose of, 76–77, 80
 - tips for, 80
- engagement metric, 172
- error messages, 105–111
 - added by engineers, 220
 - avoiding blame, 106
 - blocking errors, 109
 - categories of errors, 106
 - clarity of, 199
 - detour errors, 108

- empathy, 106
- inline errors, 106–107
- patterns, 111
- purpose of, 105, 111
- tips for, 111

- escaping characters, 221
- experience life cycle, 238–239

F

- familiarity of words and phrases, 194
- feedback from users, analyzing, 177
- Fenton, Nicole, 24
- fine-tuning LLMs, 148–149
- Fires and Foundations phase of
 - 30/60/90-day plan, 240–245
 - content strategy documentation, 244–245
 - minimum viable process, 242–244
 - tracking system, 242
- first-run experience, 171–172
- Freely, Chris, 18
- Frontitude CMS, 222
- F-shape scanning behavior, 71

G

- GenderMag, 198
- glossaries, 229–230
- goals, 1–15
 - aligning between people and organization, 2–5
 - choosing content to meet, 5–11
 - of editing, 156–157
 - identifying purposes, opportunities, and constraints, 11–12
 - imagining and testing solutions, 12–14
- Goodhart's law, 167
- Google Docs, 217
- grammar, in voice charts, 32–33
- Grammarly, 232

H

- harmlessness of content, 129
- helpfulness of content, 127
- heuristics, 182–208
 - accessibility, 187–191
 - clarity, 196–200
 - concision, 193–195

- conversationality, 194–196
- purpose, 191–192
- voice, 200–202
- hierarchical decisions, 44–45
- how-to and policy information, 8, 198–199
- HubSpot, 18
- Hurston, Zora Neale, 167

I

- idioms, 163
- important versus urgent work, 224–226
- inline errors, 106–107
- input for LLMs
 - eliciting accurate, 133–135
 - preparing test input, 139–143
- interviewing, 177–179
- intrinsic content, 7
- iterative writing, 14

J

- Jira, 222

K

- keywords, placement of, 71
- King, Stephen, 209
- K variable, 150

L

- labels, 81–85
 - avoiding complexity, 81–82
 - descriptions versus, 81
 - internationalization, 82, 84
 - patterns, 85
 - providing context through, 83
 - purpose of, 81, 85
 - tips for, 85
- languages
 - accessibility and, 187–188
 - labels, 84
- Lee, Kate Kiefer, 24
- Leonard, Elmore, 155
- LitMop
 - accuracy of content, 129, 134–135
 - adjusting model variables, 150
 - auditability of content, 131
 - buttons, 65–66
 - concepts, 28–29
 - defined, 20–21

- descriptions, 75
- design components, 228
- embedded AI, 114–115
- empty states, 79–80
- engineering prompts, 135–137
- error messages, 110
- fine-tuning, 149
- good success reports, 121–126
- grammar, 33
- harmlessness of content, 130
- helpfulness of content, 128
- human alignment, 151–152
- identifying named entities, 118
- improving, 152
- integrating into shippable product, 151–152
- labels, 84
- LLM-generated content, 119–121
- notifications, 104–105, 229
- preparing test inputs, 141–143
- product principles, 25–26
- punctuation and capitalization, 36
- purpose of, 25
- qualitative measures, 144
- sustainability, 132
- text input fields, 91–92
- titles, 61–62
- transitional text, 95–97
- verbosity, 31–32
- vocabulary, 29–30
- voice chart, 39
- LLMs (large language models), 113–154
 - accuracy of, 128–129, 133–135
 - adapting and aligning results of, 148–154
 - auditability of, 130
 - business and ethical requirements, 127–133
 - collaborative design, 51–53
 - consistency of voice, 18
 - defining “good”, 121–127
 - defining the problem, 119–121
 - designing for prompts and completions, 133–144
 - eliciting content from, 215
 - embedded AI, 114–115
 - harmlessness of, 129
 - helpfulness of, 127
 - measurement of, 143–147

- overview of, 115–118
- regulatory frameworks, 127
- selecting, 118–119
- specialization of, 118
- steps after shipping, 152–153
- sustainability, 131–133
- training, 41, 116
- transformer models, 116
- turning to for inspiration, 13
- voice charts and, 41

lo-fi tools, 211–214

M

magic wand exercise, 180–181

manifesto of content design, 251–253

MAU (monthly active users), 172

measurements for LLMs

- perplexity, 147
- qualitative, 144–145
- quantitative, 145–146
- ROGUE metrics, 146

meetings, 234–236

menus, 64–69

- brevity of items in, 64, 67
- empty states, 77
- patterns, 69
- purpose of, 69
- tips for, 69

metaphors, 163

metrics, for behaviors, 169–176

- completion, 174
- cost reduction, 175–176
- engagement, 172
- onboarding pace, 171–172
- referrals, 175
- retention, 173–174

Microsoft Tay chatbot, 133

Microsoft Word, 217

Microsoft Writing Style Guide, 231

Microsoft Xbox, 17–18

minimum viable process, 242–244

MLA Handbook, 231

monthly active users (MAU), 172

N

name design and selection, 230–231

Nicely Said (Fenton and Lee), 24

Nielsen Norman Group, 182

notifications, 101–105

- bad news, 103
- clarity, 164
- concision, 160–162
- conversationality, 162–163
- driving engagement, 103
- enticing and interesting, 102
- patterns, 105
- purposefulness, 157–160
- purpose of, 101, 105
- tips for, 105
- title and description, 101–102

O

onboarding pace, 171–172

one-shot prompts, 138–139

opportunities, identifying, 11–12

order and sequence of information, 195

organization, aligning goals of people with, 2–5

P

patterns. *See* design patterns; text patterns

perplexity, 147

personality definition, 227–228

precision metric, 146

prioritizing tasks, 224–226, 249–254

- advocating for content to the team, 250–253
- content design manifesto, 251–253
- grounding content in empathy, 249–250

processes. *See* tools and processes

product owners, proposing changes to, 167–168

product principles, 23–26, 227

- concepts supporting, 26–28
- grammar supporting, 32–33
- punctuation and capitalization supporting, 34–36
- verbosity supporting, 30–32
- vocabulary supporting, 28–30

programming knowledge, 221

prompts

- eliciting accurate input, 133–135
- engineering, 135–140
- preparing test inputs, 139–143

- publishing text, 219–222
 - content management systems, 222
 - reviewing text in code, 220–221
- punctuation, in voice charts, 34–36
- purpose
 - editing and, 157–159
 - heuristics for, 191–192
 - identifying, 11–12
- P value, 150

Q

- qualitative measures, 144–145
- quantitative measures, 145–146

R

- RAG (retrieval-augmented generation), 141–142
- ratings, 8
- reading-level accessibility, 188–189
- recall metric, 146
- Recall-Oriented Understudy for Gisting Evaluation (ROUGE) metrics, 146
- referrals, 175
- regulated content, documenting, 217–219
- regulation of LLMs, 127–133
- relationship building, 234–237
- relevance, 193
- researching content effectiveness, 167–208
 - direct behavior measurement, 169–176
 - heuristics, 182–208
 - methods for, 176–181
- research participants, recruiting, 178
- retention metric, 173–174
- retrieval-augmented generation (RAG), 141–142
- review documents, 217–219
- reviewing
 - content, 216–217
 - text in code, 220–221
- ROUGE (Recall-Oriented Understudy for Gisting Evaluation) metrics, 146

S

- scorecard for heuristic analysis, 182–183
 - example of, 185–206
 - template, 184–186
- screen readers, 190
- screenshots, 211–214
- single-task titles, 62–64, 67
- solutions, imagining and testing, 12–14
- Speth, Ralf, 1
- static content, 7
- Sturgeon Club
 - buttons, 64–65
 - concepts, 26–27
 - confirmation messages, 98–99
 - controls, 86–87
 - defined, 20–21
 - descriptions, 72–73
 - empty states, 77
 - error messages, 107
 - grammar, 32
 - labels, 81–82
 - notifications, 102
 - product principles, 24–25
 - punctuation and capitalization, 34–35
 - purpose of, 24
 - sign-in screen, 40
 - text input fields, 89–90
 - titles, 59–60
 - transitional text, 93–94
 - verbosity, 30
 - vocabulary, 28–29
 - voice chart, 37
- style guides, 231
- surveys, 180–181
- sustainability of LLMs, 131–133
- Swift, 221
- systematization of content design, 226–232
 - patterns, 228–231
 - personality, 227–228
 - practicalities, 231–232
 - principles, 227

T

- TAPP
 - buttons, 67–68
 - clarity, 164
 - concepts, 27

- concision, 160–162
- confirmation messages, 100
- controls, 86–88
- conversational design, 49–55
- conversationality, 162–163
- defined, 20–21
- descriptions, 73–74
- empty states, 78–79
- error messages, 109–110
- experience life cycle, 238
- goals, 2–7
- grammar, 33
- notifications, 103–104
- product principles, 25
- punctuation and capitalization, 36–37
- purposefulness, 157–160
- purpose of, 25
- sign-in screen, 40
- text input fields, 91–92
- titles, 62–63
- transitional text, 95–96
- verbosity, 31
- vocabulary, 29
- voice chart, 39
- Tay chatbot, 133
- temperature, 150
- template for content scoring, 184–186
- “10 Usability Heuristics for User Interface Design” (Nielsen Norman Group), 182
- terminology
 - consistency of, 200
 - systematizing, 229–230
- text input fields, 89–94
 - hint text and labels, 89–90
 - patterns, 93
 - prefilling, 89
 - purpose of, 89, 93
 - tips for, 93
- text patterns, 57–111
 - buttons, 64–69
 - confirmation messages, 98–101
 - controls, 85–88
 - descriptions, 69–76
 - empty states, 76–81
 - error messages, 105–111
 - labels, 81–85
 - menus, 64–69
 - notifications, 101–105
 - text input fields, 89–94
 - titles, 58–64
 - transitional text, 93–98
 - versus LLMs, 215–216
- 30/60/90-day plan, 233–247
 - days 0–30, 233–240
 - days 31–60, 240–245
 - days 61–90, 245–246
- tie breaking, 43–44
- titles, 58–64
 - brevity of, 59
 - context-setting, 59–62
 - matching wording in buttons, 67
 - patterns, 63
 - purpose of, 58, 63
 - single-task, 62–64
 - tips for, 63
- tokens, 116
- tone, defined, 23
- tools and processes
 - documenting regulated content, 217–219
 - managing content review, 216–217
 - prioritizing urgent and important tasks, 224–226
 - publishing text, 219–222
 - surveying, 180
 - systemization of content design, 226–232
 - tracking content work, 222–224
 - writing for context, 209–216
- tool tips, 75
- tracking systems, 222–224, 242
- transformer models, 116
- transitional text, 93–98
 - heightening excitement, 94
 - patterns, 96
 - purpose of, 93, 96
 - specificity of, 95
 - tips for, 96
- troubleshooting content, 8

U

- units, 231–232
- urgent versus important work, 224–226
- usability
 - heuristics scorecard, 183–185
 - testing, 179
 - voice and, 202

UX (user experiences)
30/60/90-day plan, 233–247
content-first design, 47–56
editing, 155–165
goals, 1–15
LLM-based experiences, 113–154
prioritizing content design
work, 249–254
researching content effective-
ness, 167–208
strategic purpose of content
design, 2
tools and processes, 209–232
voice of, 17–45
UX writers. *See* content designers

V

vectors, 116
verbosity, in voice charts, 30–32
virtuous cycle
aligning goals, 2–4
choosing content to meet
goals, 5–10
experience life cycle, 238
vocabulary, in voice charts, 28–30
voice, 17–45
consistency of, 17–18
example experiences, 19–22
heuristics for, 200–202
heuristics scorecard, 183–185
role in retention, 173
voice charts
completing, 37–41
concepts, 26–28
content design systems, 227
creating, 22–40

decision making and tie break-
ing, 43–44
designing new content, 42–43
grammar, 32–33
informing LLM training, 41
product principles, 23–26
punctuation and capitaliza-
tion, 34–36
roles of, 22
tone, 23
training new content design-
ers, 41
using, 41–46
verbosity, 30–32
vocabulary, 28–30

W

WCAG (Web Content Accessibility
Guidelines), 182
What and Who phase of 30/60/90-
day plan, 233–240
experience life cycle, 238–239
relationship building, 234–237
starting to write, 239–240
working groups, 12–13
work items, 222–224
Writer.ai, 232

X

Xbox, 17–18
XML, 221

Z

zero-shot prompts, 137, 139

About the Author

Torrey Podmajersky helps teams solve business and customer problems using UX content. She has written inclusive and accessible experiences for consumer, technical, and business audiences in Google Ads, Google Cloud, Google Support, OfferUp, Xbox, and Microsoft. Torrey consults with B2B, B2C, and mentoring clients, and teaches and speaks worldwide about UX and content design through her agency, Catbird Content.

Colophon

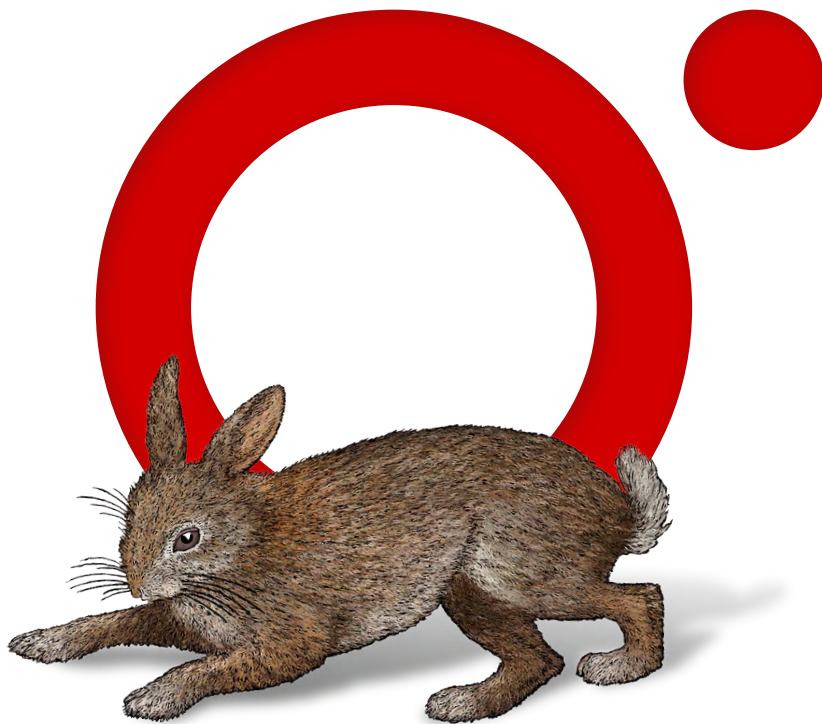
The animal on the cover of *Strategic Writing for UX* is a gray catbird (*Dumetella carolinensis*). This group of birds has a wailing call that resembles the meow of a cat, giving them their common name. They are found in North and Central America, favoring a habitat of scrubland and the outskirts of forests.

The gray catbird has mostly gray plumage, with rust-colored coverts under its tail. Though the coloring of this species is nondescript, they have vivid personalities: apart from the cat-like noise, they have a variety of calls (including imitations of other birds) and quirky physical motions like drooping the wings, angling the head, and fanning out the tail. Male catbirds often “riff” for over 10 minutes, at a rate of 90 syllables per minute.

Catbirds have a diet of insects and berries, which they forage for on the ground or within shrubs and trees. Females build cup-shaped nests about 3–10 feet off the ground, and lay 2 or 3 light-blue eggs at a time. The catbird is adept at distinguishing its own eggs from those of the brown cowbird (who opportunistically lays eggs in other birds’

nests so that the other mother will mistake the cowbird's chicks for her own and raise them). Catbirds will instead eject the incorrect eggs from the nest.

The cover image is a color illustration by Karen Montgomery, based on a black-and-white engraving from *Brehms Thierleben*. The series design is by Edie Freedman, Ellie Volckhausen, and Karen Montgomery. The cover fonts are Gilroy Semibold and Guardian Sans. The text font is Scala; and the heading font is Gotham.



O'REILLY®

**Learn from experts.
Become one yourself.**

60,000+ titles | Live events with experts
Role-based courses | Interactive learning
Certification preparation



**Try the O'Reilly learning platform
free for 10 days.**