Francesco Corea

An Introduction to Data

Everything You Need to Know About AI, Big Data and Data Science

Springer

Studies in Big Data

Volume 50

The series "Studies in Big Data" (SBD) publishes new developments and advances in the various areas of Big Data—quickly and with a high quality. The intent is to cover the theory, research, development, and applications of Big Data, as embedded in the fields of engineering, computer science, physics, economics and life sciences. The books of the series refer to the analysis and understanding of large, complex, and/or distributed data sets generated from recent digital sources coming from sensors or other physical instruments as well as simulations, crowd sourcing, social networks or other internet transactions, such as emails or video click streams and other. The series contains monographs, lecture notes and edited volumes in Big Data spanning the areas of computational intelligence incl. neural networks, evolutionary computation, soft computing, fuzzy systems, as well as artificial intelligence, data mining, modern statistics and Operations research, as well as self-organizing systems. Of particular value to both the contributors and the readership are the short publication timeframe and the world-wide distribution, which enable both wide and rapid dissemination of research output.

More information about this series at http://www.springer.com/series/11970

Francesco Corea

An Introduction to Data

Everything You Need to Know About AI,
Big Data and Data Science

Francesco Corea
Department of Management
Ca' Foscari University
Venice, Italy

To you, now and forever

Preface

This book aims to be an introduction to big data, artificial intelligence and data science for anyone who wants to learn more about those domains. It is neither a fully technical book nor a strategic manual, but rather a collection of essays and lessons learned doing this job for a while.

In that sense, this book is not an organic text that should be read from the first page onwards, but rather a collection of articles that can be read at will (or at need). The structure of the chapter is very similar, so I hope the reader won't find difficulties in establishing comparisons or understanding the differences between specific problems AI is being used for. I personally recommend reading the first three-four chapters in a row to have a general overview of the technologies and then jump around depending on what topic interests you the most.

The book also replicates some of the contents already introduced in previous research as well as shows new material created working as a data scientist, as a startup advisor as an investor. It is therefore to some extent both a new book and a 2.0 version of some previous work of mine, but for sure the content is reorganized in a completely new way and gets new meaning when read in a different context.

Artificial intelligence is certainly a hot topic nowadays, and this book wants to be both a guide on the past and a tool to look into the future. I always tried to maintain a balance between explaining concepts, tools and ways in which AI has been used, and *potential* applications or trends for future. I hope the reader may find himself not only grasping how relevant AI, big data and data science are for our progress as a society, but also wondering *what's next*.

The book is structured in such a way that the first few chapters explain the most relevant definitions and business contexts where AI and big data can have an impact. The rest of the book looks instead at specific sectorial applications, issues or more generally subjects that AI is meaningfully changing.

Finally, I am writing this book hoping that it will be valuable for some readers in how they think and use technologies to improve our lives and that it could stimulate conversations or projects that could produce a positive impact in our society.

Venice, Italy   Francesco Corea

Contents

List of Figures

List of Tables

Chapter 1
Introduction to Data

There are many ways to define what big data is, and this is probably why it still remains a really difficult concept to grasp. Today, someone describes big data as dataset above a certain threshold, e.g., over a terabyte (Driscoll 2010), others as data that crash conventional analytical tools like Microsoft Excel. More renowned works though identified big data as data that display features of *Variety*, *Velocity*, and *Volume* (Laney 2001; McAfee and Brynjolfsson 2012; IBM 2013; Marr 2015). Even though they are all partially true, there is a definition that seems to better capture this phenomenon (Dumbill 2013; De Mauro et al. 2015; Corea 2016): big data analytics is an innovative approach that consists of different technologies and processes to extract worthy insights from low-value data that do not fit, for any reason, the conventional database systems.

In the last few years the academic literature on big data has grown extensively (Lynch 2008). It is possible to find specific applications of big data to almost any field of research (Chen et al. 2014). For example, big data applications can be found in medical-health care (Murdoch and Detsky 2013; Li et al. 2011; Miller 2012a, b); biology (Howe et al. 2008); governmental projects and public goods (Kim et al. 2014; Morabito 2015); financial markets (Corea 2015; Corea and Cervellati 2015). In other more specific examples, big data have been used for energy control (Moeng and Melhem 2010), anomaly detection (Baah et al. 2006), crime prediction (Mayer-Schönberger and Cukier 2013), and risk management (Veldhoen and De Prins 2014).

No matter what business is considered, big data are having a strong impact on every sector: Brynjolfsson et al. (2011) proved indeed that a data-driven business performs between 5 and 6% better than its competitors. Other authors instead focused their attention on organizational and implementation issues (Wielki 2013; Mach-Król et al. 2015). Marchand and Peppard (2013) indicated five guidelines for a successful big data strategy: (i) placing people at the heart of Big Data initiatives; (ii) highlighting information utilization to unlock value; (iii) adding behavioral scientists to the team; (iv) focusing on learning; and (v) focusing more on business problems than technological ones. Barton and Court (2012) on the other hand

identified three different key features for exploiting big data potential: choosing the right data, focusing on biggest driver of performance to optimize the business, and transforming the company's capabilities.

Data are quickly becoming a new form of capital, a different coin, and an innovative source of value. It has been mentioned above how relevant it is to channel the power of the big data into an efficient strategy to manage and grow the business. But it is also true that big data strategies may not be valuable for all businesses, mainly because of structural features of the business/company itself. However, it is certain that a data strategy is still useful, no matter the size of your data. Hence, in order to establish a data framework for a company, there are first of all few misconceptions that need to be clarified:

i) **More data means higher accuracy**. Not all data are good quality data, and tainting a dataset with dirty data could compromise the final products. It is similar to a blood transfusion: if a non-compatible blood type is used, the outcome can be catastrophic for the whole body. Secondly, there is always the risk of overfitting data into the model, yet not derive any further insight—"if you torture the data enough, nature will always confess" (Coase 2012). In all applications of big data, you want to avoid striving for perfection: too many variables increase the complexity of the model without necessarily increasing accuracy or efficiency. More data always implies higher costs and not necessarily higher accuracy. Costs include: higher maintenance costs, both for the physical storage and for model retention; greater difficulties in calling the shots and interpreting the results; more burdensome data collection and time-opportunity costs. Undoubtedly the data used do not have to be orthodox or used in a standard way—and this is where the real gain is locked in—and they may challenge the conventional wisdom, but they have to be proven and validated. In summary, smart data strategies always start from analyzing internal datasets, before integrating them with public or external sources. Do not store and process data just for data's sake, because with the amount of data being generated daily, the noise increases faster than the signal (Silver 2013). Pareto's 80/20 rule applies: the 80% of the phenomenon could be probably explained by the 20% of the data owned.

ii) **If you want to do big data**, **you have to start big**. A good practice before investing heavily in technology and infrastructures for big data is to start with few high-value problems that validate whether big data may be of any value to your organization. Once the proof of concept demonstrates the impact of big data, the process can be scaled up.

iii) **Data equals Objectivity**. First of all, data need to be contextualized, and their "objective" meaning changes depending on the context. Even though it may sound a bit controversial, data can be perceived as objective—when it captures facts from natural phenomena—or subjective—if it reflects pure human or social constructs. In other words, data can be *factual*, i.e., the ones that are univocally the same no matter who is looking at them, or *conventional/social* —the more abstract data, which earn its right to representativeness from the

general consensus. Think about this second class of data as the notions of value, price, and so forth. It is important to bear this distinction in mind because the latter class is easier to manipulate or can be victim of a self-fulfilling prophecy. As stated earlier on, the interpretation of data is the quintessence of its value to business. Ultimately, both types of data could provide different insights to different observers due to relative problem frameworks or interpretation abilities (the so-called *framing effect*). Data science will therefore never be a proper science, because it will lack of full objectivity and full replicability, and because not every variable can be precisely quantified, but only approximated.

Let's also not forget that a wide range of behavioral biases that may invalidate the objectivity of the analysis affects people. The most common ones between both scientists and managers are: *apophenia* (distinguishing patterns where there are not), *narrative fallacy* (the need to t patterns to series of disconnected facts), *confirmation bias* (the tendency to use only information that confirms some priors)—and his corollary according to which the search for evidences will eventually end up with evidences discovery—and *selection bias* (the propensity to use always some type of data, possibly those that are best known). A final interesting big data curse to be pointed out is nowadays getting known as the "Hathaway's effect": it looked like that where the famous actress appeared positively in the news, Warren Buffett's Berkshire Hathaway company observed an increase in his stock price. This suggests that sometime there exist correlations that are either spurious or completely meaningless and groundless.

iv) **Your data will reveal you all the truth**. Data on its own are meaningless, if you do not pose the right questions first. Readapting what DeepThought says in *The Hitchhikers' Guide to the Galaxy* written by Adams many years ago, big data can provide the final answer to life, the universe, and everything, as soon as the right question is asked. This is where human judgment comes into: posing the right question and interpreting the results are still competence of the human brain, even if a precise quantitative question could be more efficiently replied by any machine.

The alternative approach that implements a random data discovery—the so-called "*let the data speak*" approach—is highly inefficient, resource consuming and potentially value-destructive. An intelligent data discovery process and exploratory analysis therefore is highly valuable, because "*we don't know what we don't know*" (Carter 2011).

The main reasons why data mining is often ineffective is that it is undertaken without any rationale, and this leads to common mistakes such as false positives, overfitting, ignoring spurious relations, sampling biases, causation-correlation reversal, wrong variables inclusion or model selection (Doornik and Hendry 2015; Harford 2014). A particular attention has to be put on the causation-correlation problem, since observational data only take into account the second aspect. However, According to Varian (2013) the problem can be solved through experimentations.

References

Baah, G. K., Gray, A., Harrold, M. J. (2006). Online anomaly detection of deployed software: A statistical machine learning approach. In *Proceedings of the 3rd International Workshop on Software Quality Assurance*, 70–77.

Barton, D., & Court, D. (2012). Making advanced analytics work for you. *Harvard Business Review, 90*(10), 78–83.

Brynjolfsson, E., Hitt, L. M., & Kim, H. H. (2011). Strength in numbers: How does data-driven decision making affect firm performance?. Available at SSRN: http://ssrn.com/abstract= 1819486.

Carter, P. (2011). Big data analytics: Future architectures, skills and roadmaps for the CIO. *IDC White Pape*r, Retrieved from http://www.sas.com/resources/asset/BigDataAnalytics-FutureArchitectures-Skills-RoadmapsfortheCIO.pdf.

Chen, M., Mao, S., Zhang, Y., Leung, V.C. (2014). Big data: Related technologies, challenges and future prospects, SpringerBriefs in Computer Science, 59.

Coase, R. H. (2012). *Essays on economics and economists*, University of Chicago Press.

Corea, F. (2015). Why social media matters: the use of Twitter in portfolio strategies. *International Journal of Computer Applications, 128*(6), 25–30.

Corea, F. (2016). *Big data analytics: A management perspective. Studies Series in Big Data*, 21. Springer International Publishing.

Corea, F., & Cervellati, E. M. (2015). The power of micro-blogging: How to use twitter for predicting the stock market. *Eurasian Journal of Economics and Finance, 3*(4), 1–6.

De Mauro, A., Greco, M., & Grimaldi, M. (2015). What is big data? A consensual definition and a review of key research topics. *AIP Conference Proceedings, 1644,* 97–104.

Doornik, J. A., & Hendry, D. F. (2015). Statistical model selection with big data. *Cogent Economics & Finance, 3,* 1045216.

Driscoll, M. E. (2010, Dec. 20). How much data is "Big Data"?, [Msg 2]. Message posted to https://www.quora.com/How-much-data-is-Big-Data.

Dumbill, E. (2013). Making sense of big data. *Big Data, 1*(1), 1–2.

Harford, T. (2014). Big data: Are we making a big mistake? Financial Times. Retrieved fromhttp://www.ft.com/cms/s/2/21a6e7d8-b479-11e3-a09a-00144feabdc0.html#ixzz2xcdlP1zZ.

Howe, A. D., Costanzo, M., Fey, P., Gojobori, T., Hannick, L., Hide, W. … Rhee, S. Y. (2008). Big data: The future of biocuration. *Nature*, *455*(7209), 47–50.

IBM (2013). The four V's of big data. Retrieved from http://www.ibmbigdatahub.com/infographic/four-vs-big-data.

Kim, G. H., Trimi, S., & Chung, J. H. (2014). Big-data applications in the government sector. *Communications of the ACM, 57*(3), 78–85.

Laney, D. (2001). 3D data management: Controlling data volume, velocity, and variety. META group Inc. https://blogs.gartner.com/doug-laney/files/2012/01/ad949-3D-Data-Management-Controlling-Data-Volume-Velocity-and-Variety.pdf. Accessed on Oct 27, 2015.

Li, Y., Hu, X., Lin, H., & Yang, Z. (2011). A framework for semisupervised feature generation and its applications in biomedical literature mining. *IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB), 8*(2), 294–307.

Lynch, C. (2008). Big data: How do your data grow? *Nature, 455,* 28–29.

Mach-Król, M., Olszak, C. M., Pełech-Pilichowski, T. (2015). *Advances in ICT for business, industry and public sector. Studies in Computational Intelligence*. Springer, 200 pages.

Marchand, D., & Peppard, J. (2013). Why IT fumbles analytics. *Harvard Business Review, 91*(1/2), 104–113.

Marr, B. (2015). *Big data: Using SMART big data, analytics and metrics to make better decisions and improve performance*. Wiley, 256 pages.

Mayer-Schönberger, V., Cukier, K. (2013). Big data: A revolution that will transform how we live, work, and think*. Eamon Dolan/Houghton Mifflin Harcourt.

McAfee, A., & Brynjolfsson, E. (2012). Big data: The management revolution. *Harvard Business Review, 90*(10), 60–66.

Miller, K. (2012a). Leveraging social media for biomedical research: How social media sites are rapidly doing unique research on large cohorts. *Biomedical Computation Review* (available at http://biomedicalcomputationreview.org/content/leveraging-social-media-biomedical-research. Accessed October 27, 2015).

Miller, K. (2012b). Big data analytics in biomedical research. *Biomedical Computation Review* (available at https://biomedicalcomputationreview.org/content/big-data-analytics-biomedical-research. Accessed October 27, 2015).

Moeng, M., Melhem, R. (2010). Applying statistical machine learning to multicore voltage and frequency scaling. In *Proceedings of the 7th ACM international conference on Computing frontiers*, 277–286.

Morabito, V. (2015). *Big data and analytics: Strategic and organizational impacts*. Springer International Publishing, 183 pages.

Murdoch, T. B., & Detsky, A. S. (2013). The inevitable application of big data to health care. *JAMA, 309*(13), 1351–1352.

Silver, N. (2013). *The signal and the noise: The art and science of prediction*. Penguin.

Varian, H. (2013). Beyond big data. NABE annual meeting, San Francisco, CA, September 10th, 2013.

Veldhoen, A., De Prins, S. (2014). Applying big data to risk management. *Avantage Reply Report*, 1–14.

Wielki, J. (2013). Implementation of the Big Data concept in organizations—possibilities, impediments, and challenges. In *Proceedings of the 2013 Federated Conference on Computer Science and Information Systems*, 985– 989.

Chapter 2
Big Data Management: How Organizations Create and Implement Data Strategies

It has been mentioned above how relevant it is to channel the power of the big data into an effective strategy to manage and grow the business. However, a consensus on how and what to implement is difficult to be achieved and what is then proposed is only one possible approach to the problem.

Following the guidelines given by Doornik and Hendry (2015), we find a lean approach to data problem to be not only useful but above all efficient. It actually reduces time, effort and costs associated with data collection, analysis, technological improvements and ex-post measuring. The relevance of the framework lies in avoiding the extreme opposite situations, namely collecting all or no data at all. The Fig. 2.1 illustrates key steps towards this lean approach to big data: first of all, business processes have to be identified, followed by the analytical framework that has to be used. These two consecutive stages have feedback loops, as well as the definition of the analytical framework and the dataset construction, which has to consider all the types of data, namely data at rest (static and inactively stored in a database), at motion (inconstantly stored in temporary memory), and in use (constantly updated and store in database). The modeling phase is crucial, and it embeds the validation as well, while the process ends with the scalability implementation and the measurement. A feedback mechanism should prevent an internal stasis, feeding the business process with the outcomes of the analysis instead of improving continuously the model without any business response.

Data need to be consistently aggregated from different sources of information, and integrated with other systems and platforms; common reporting standards should be created—the master copy—and any information should need to be eventually validated to assess accuracy and completeness. Finally, assessing the skills and profiles required to extract value from data, as well as to design efficient data value chains and set the right processes, are two other essential aspects. Having a solid internal data management, jointly with a well-designed golden record, helps to solve the huge issue of *stratified entrance*: dysfunctional datasets resulting from different people augmenting the dataset at different moments or across different layers.

Define the
business
process

•Identifying
customer's
objectives;
•Frame the
problem;
•State a
business
hypothesis;
•Sample
dataset and
exploratory
analysis.

Define the
analytical
framework

•Identifying
models to be
used
•Design
implementatio
n steps

Create the full
dataset

•Data
acquisition:
•Data at rest:
sampling
•Data in
motion:
processing
•Data in use:
capture and
processing

Modelling
Phase

•Model
construction
•Model
Validation
•Model
Visualization

Optimization

•Scale
•Measurement
•Monitoring
•Calibration

Fig. 2.1 Big data lean deployment approach

Even if a data lean approach is used, companies may incur many problems. It is essential then to develop a framework to track internal developments and obstacles, as well as to draw the next steps in the analytics journey. A *Data Stage of Development Structure* (DS2) is a maturity model built for this purpose, a roadmap developed to implement a revenue-generating and impactful data strategy. It can be used to assess a company's current situation, and to understand the future steps to undertake to enhance internal big data capabilities.

Table 2.1 provides a four by four matrix where the increasing stages of evolution are indicated as *Primitive*, *Bespoke*, *Factory*, and *Scientific*, while the metrics they are considered through are *Culture*, *Data*, *Technology*, and *Talent*. The final considerations are drawn in the last row, the one that concerns the financial impact on the business of a well-set data strategy.

Stage one is about raising awareness: the realization that data science could be relevant to the company business. In this phase, there are neither any governance structures in place nor any pre-existing technology, and above all no organization-wide buy-in. Yet, tangible projects are still the result of individual's data enthusiasm being channeled into something actionable. The set of skills owned is still rudimental, and the actual use of data is quite rough. Data are used only to convey basic information to the management, so it does not really have any impact on the business. Being at this stage does not mean being inevitably unsuccessful, but it simply shows that the projects performance and output are highly variable, contingent, and not sustainable. The second Phase is the reinforcing: it is actually an exploration period. The pilot has proved big data to have a value, but new competencies, technologies, and infrastructures are required—and especially a new data governance, in order to also take track of possible data contagion and different

Table 2.1 Data stage of development structure

Drivers/ stages	Primitive	Bespoke	Factory	Scientific
Culture	• No leadership support • Analytics as an IT asset • Conveying information (reporting, dashboard, etc.) • No budget • Descriptive analytics	• Leadership interest and midlevel management backing • Analytics used to understand problems • Specific application/ department • Funding for specific project • Tailored modus operandi (not replicable) • Predictive analytics	• Leadership sponsorship • Analytics used to identify issues and develop actionable options • Alignment to the business as a whole • Specific budget for analytics function • Advanced data mining • Prescriptive analytics	• Full executive support • Data driven business • Cross-department applications • Substantial infrastructural, human, and technology investments • Advanced data discovery • Automated analytics
Data	• Absence of a proper data infrastructure • Disorganized and dispersed silos • Duplicated information	• Data marts (lack of variety) • Internal structured data points • Data gaps or incomplete	• Virtual data marts • Internal and external data, • Mainly structured data • Easy to manage unstructured data (e.g., texts)	• Data lakes • Any data (unstructured, semi-structured, etc.) • Variety of sources (IoT, Social media, etc.) • Information life cycle in place
Technology	• Absence of data governance • No forefront technology (spreadsheet for reporting) • Low investments	• Integrated relational database (SQL) • Improvements in data architecture • Setting of a golden record • Scripting languages	• Pioneering technologies (Hadoop, MapReduce—see Appendix 1) • Integration with programming languages • Visualization tools	• Centralized dataset • Cloud storage • Mobile applications • APIs, internet of things, and advanced machine learning tools
Talent	• Dispersed talents • Few people with few data analytical skills	• Mix of few full time and some part-time data scientists • Proper data warehouse team • Strategic partnership fur enhancing capabilities	• Well framed recruitment process • Proper data science team • IT department fully formed and operative • Supporting of IT to data team	• Centre of excellence • Dominion experts and specialists • Training and continuous learning • Active presence within the Data Ecosystem
Impact	No return on Investments (ROI)	Moderate revenues, that Justify though further investments	Significant revenues	Revolutionized business model (blue ocean revenues)

actors who enter the data analytics process at different stages. Since management's contribution is still very limited, the potential applications are relegated to a single department or a specific function. The methods used although more advanced than in Phase one are still highly customized and not replicable. By contrast, Phase three adopts a more standardized, optimized, and replicable process: access to the data is much broader, the tools are at the forefront, and a proper recruitment process has been set to gather talents and resources. The projects benefit from regular budget allocation, thanks to the high-level commitment of the leadership team. Step four deals with the business transformation: every function is now data-driven, it is led by agile methodologies (i.e., deliver value incrementally instead of at the end of the production cycle), and the full-support from executives is translated into a series of relevant actions. These may encompass the creation of a Centre of Excellence (i.e., a facility made by top-tier scientists, with the goal of leveraging and fostering research, training and technology development in the field), high budget and levels of freedom in choosing the scope, or optimized cutting-edge technological and architectural infrastructures, and all these bring a real impact on the revenues' flow. A particular attention has to be especially put on data lakes, repositories that store data in native formats: they are low costs storage alternatives, which supports manifold languages. Highly scalable and centralized stored, they allow the company to switch without extra costs between different platforms, as well as guarantee a lower data loss likelihood. Nevertheless, they require a metadata management that contextualizes the data, and strict policies have to be established in order to safeguard the data quality, analysis, and security. Data must be correctly stored, studied through the most suitable means, and to be breach-proof. An information lifecycle has to be established and followed, and it has to take particular care of timely efficient archiving, data retention, and testing data for the production environment.

A final consideration has to be spared about cross-stage dimension "culture". Each stage has associated a different kind of analytics, as explained in Davenport (2015). Descriptive analytics concerned what happened, predictive analytics is about future scenarios (sometimes augmented by diagnostic analytics, which investigates also the causes of a certain phenomenon), prescriptive analytics suggests recommendations, and finally, automated analytics are the ones that take action automatically based on the analysis' results.

Some of the outcomes presented so far are summarized in Fig. 2.2. The following chart shows indeed the relationship between management's support for the analytics function and the complexity and skills required to excel into data- driven businesses. The horizontal axis shows the level of commitment by the management (high vs. low), while the vertical axis takes into account the feasibility of the project undertaken—where feasibility is here intended as the ratio of the project's complexity and the capabilities needed to complete it. The intersection between feasibility of big data analytics and management involvement divides the matrix into four quarters, corresponding to the four types of analytics. Each circle identifies one of the four stages (numbered in sequence, from I—*Primitive*, to IV—*Scientific*). The size of each circle indicates its impact on the business (i.e., the larger the circle, the higher the ROI). Finally, the second horizontal axis keeps track of the increasing
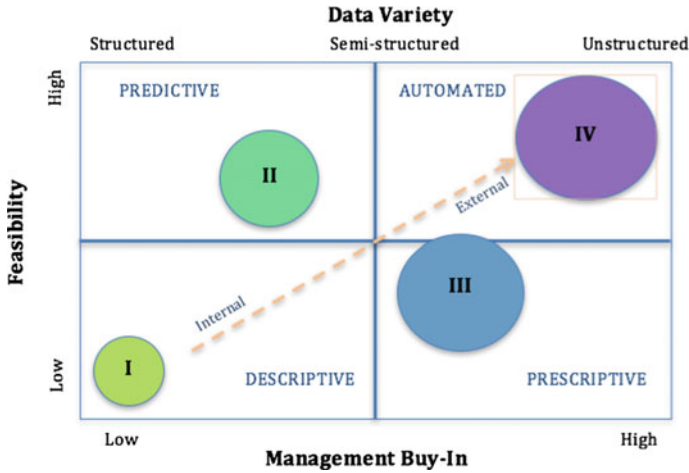
Fig. 2.2 Big data maturity map

data variety used in the different stages, meaning structure, semi-structured, or unstructured data (i.e., IoT, sensors, etc.). The orange diagonal represents what kind of data are used: from closed systems of internal private networks in the bottom left quadrant to market/public and external data in the top right corner.

Once the different possibilities and measurements have been identified (see the Appendix II or Corea 2016 for the full details on the framework), they can be used to understand what stage a firm belongs to. It is also useful to see how a company could transition from one level to the next and in the following figure some recommended procedures have been indicated to foster this transition.

In order to smoothly move from the *Primitive* stage to the *Bespoke*, it is necessary to proceed by experiments run from single individuals, who aim to create proof of concepts or pilots to answer a single small question using internal data. If these questions have a good/high-value impact on the business, they could be acknowledged faster. Try to keep the monetary costs low as possible (cloud, open source, etc.), since the project will be already expensive in terms of time and manual effort. On a company level, the problem of data duplication should be addressed. The transition from *Bespoke* to *Factory* instead demands the creation of standard procedures and golden records, and a robust project management support. The technologies, tools, and architecture have to be experimented, and thought as they are implemented or developed to stay. The vision should be medium/long term then. It is essential to foster the engagement of the higher- senior management level. At a higher level, new sources and type of data have to be promoted, data gaps have to be addressed, and a strategy for platforms resiliency should be developed. In particular, it has to be drawn down the acceptable data loss (*Recovery Point Objective*), and the expected recovered time for disrupted units (*Recovery Time Objective*). Finally, to become data experts and leaders and shifting to the *Scientific* level, it is indispensable to focus on details, optimize models and datasets, improve

the data discovery process, increase the data quality and transferability, and identifying a blue ocean strategy to pursue. Data security and privacy are essential, and additional transparency on the data approach should be released to the shareholders. A Centre of Excellence (CoE) and a talent recruitment value chain play a crucial role as well, with the final goal to put the data science team in charge of driving the business. The CoE is indeed fundamental in order to mitigate the short-term performance goals that managers have, but it has to be reintegrated at some point for the sake of scalability. It would be possible now to start documenting and keeping track of improvements and ROI. From the final step on, a process of continuous learning and forefront experimentations is required to maintain a leadership and attain respectability in the data community.

In Fig. 2.3 it has also been indicated a suggested timeline for each step, respectively up to six months for assessing the current situation, doing some research and starting a pilot; up to one year for exploiting a specific project to understand the skills gap, justify a higher budget allocations, and plan the team expansion; two to four years to verify the complete support from every function and level within the firm, and finally at least five years to achieving a fully-operationally data-driven business. Of course, the time needed by each company varies due to several factors, so it should be highly customizable.

A few more words should be spent regarding the organizational home for data analytics (Pearson and Wegener 2013). We claimed that the Centre of Excellence is the cutting-edge structure to incorporate and supervise the data functions within a company. Its main task is to coordinate cross-units' activities, which embeds: maintaining and upgrading the technological infrastructures; deciding what data have to be gathered and from which department; helping with the talents recruitment; planning the insights generation phase, and stating the privacy, compliance, and ethics policies. However, other forms may exist, and it is essential to know them since sometimes they may fit better into the preexisting business model.

Figure 2.4 shows different combinations of data analytics independence and business models. It ranges between business units (BUs) that are completely independent one from the other, to independent BUs that join the efforts in some specific projects, to an internal (corporate center) or external (center of excellence) center that coordinates different initiatives.

Fig. 2.3 Maturity stage transitions

Fig. 2.4 Data analytics organizational models

In spite of everything, all the considerations made so far mean different things and provide singular insights depending on the firm's peculiarities. In particular, the different business life cycle phase in which the company is operating deeply influences the type of strategy to be followed, and it is completely unrelated to the maturity data stage to which they belong (i.e., a few months old company could be a *Scientific* firm, while a big investment bank only a *Primitive* one).

References

Corea, F. (2016). *Big Data Analytics: A Management Perspective. Studies Series in Big Data*, *21*. Springer International Publishing.

Davenport, T. H. (2015). The rise of automated analytics. *The Wall Street Journal*, January 14th 2015. Accessed on October 30th 2015 (available at http://www.tomdavenport.com/wp-content/uploads/The-Rise-of-Automated-Analytics.pdf).

Doornik, J. A., & Hendry, D. F. (2015). Statistical model selection with big data. *Cogent Economics & Finance, 3,* 1045216.

Pearson, T., Wegener, R. (2013). Big data: the organizational challenge, Bain & Company White paper.

Chapter 3
Introduction to Artificial Intelligence

Artificial Intelligence (AI) represents nowadays a paradigm shift that is driving at the same time the scientific progress as well as the industry evolution. Given the intense level of domain knowledge required to really appreciate the technicalities of the artificial engines, what AI is and can do is often misunderstood: the general audience is fascinated by its development and frightened by terminator-like scenarios; investors are mobilizing huge amounts of capital but they have not a clear picture of the competitive drivers that characterize companies and products; and managers are rushing to get their hands on the last software that may improve their productivities and revenues, and eventually their bonuses.

Even though the general optimism around creating advancements in artificial intelligence is evident (Muller and Bostrom 2016), in order to foster the pace of growth facilitated by AI I believe it would be necessary to clarify some concepts. The intent of this work is then manifold: explaining and defining few relevant terms, summarizing history of AI as well as literature advancements; investigating further innovation that AI is bringing both in scientific and business models terms; understanding where the value lies for investors; and eventually stimulating discussion about risk and future developments driven by AI.

3.1 Basic Definitions and Categorization

First, let's describe what artificial intelligence means. According to Bostrom (2014), AI today is perceived in three different ways: it is something that might answer all your questions, with an increasing degree of accuracy ("the Oracle"); it could do anything it is commanded to do ("the Genie"), or it might act autonomously to pursue a certain long-term goal ("the Sovereign"). However, AI should not be defined by what it can do or not, and thus a broader definition is appropriate.

An artificial intelligence is a system that can learn how to learn, or in other words a series of instructions (an algorithm) that allows computers to write their own algorithms without being explicitly programmed for.

Although we usually think about intelligence as the computational part of our ability to achieve certain goals, it is rather the capacity to learn and solve new problems in a changing environment. In a primordial world then, it is simply the attitude to foster survival and reproduction (Lo 2012, 2013; Brennan and Lo 2011, 2012). A living being is then defined as intelligent if she is driving the world into states she is optimizing for.

No matter how accurately we defined this concept, we can intuitively understand that the level of intelligence machines are provided with today is years far from the average level of any human being. While human being actions proceed from observing the physical world and deriving underlying relationships that link cause and effect in natural phenomena, an artificial intelligence is moved entirely by data and has no prior knowledge of the nature of the relationship among those data. It is then "artificial" in this sense because it does not stem from the physical law but rather from pure data.

We then have just defined what artificial intelligence is and what mean to us. In addition to that, though, there are two other concepts that should be treated as part of this introduction to AI: first of all, how AI is different and/or related to other buzzwords (big data, machine learning, etc.); second, what features a system has to own to be defined intelligent.

I think of AI as an interdisciplinary field, which covers (and requires) the study of manifold sub-disciplines, such as natural language processes, computer vision, as well as Internet of things and robotics. Hence, in this respect, AI is an umbrella term that gathers a bucket of different aspects. We can somehow look at AI to be similar to a fully-functional living being, and we can establish comparisons to figure out the degree of relationship between AI and other (sub)fields. If AI and the human body are alike, it has to possess a brain, which carries out a variety of tasks and is in charge of specific functions such the language (NLP), the sight (computer vision), and so on so forth. The body is made of bones and muscles, as much as a robot is made by circuits and metals. Machine learning can be seen as specific movements, action or thoughts we develop and that we fine-tune by doing. The Internet of things (IoT) corresponds to the human senses, which is the way in which we perceive the world around us. Finally, big data is the equivalent of the food we eat and the air we breathe, i.e., the fuel that makes us tick, as well as every input we receive from the external world that is captured by our senses. It is a quite rough comparison, but it conveys a simple way on how all the terms are related to each other.

Although many other comparisons may be done, and many of them can be correct simultaneously, the choice of what kind of features a system should have to be a proper AI is still quite controversial. In my opinion, the system should be endowed with a learning structure, an interactive communication interface, and a sensorial-like input digestion. Unfortunately, this idea is not rigorous from a scientific point of view, because it would involve a series of ethical, psychological, and philosophical considerations that should be taken into account.

Instead of focusing much longer on this not-provable concept, I rather prefer to illustrate how those characteristics would reflect the different types of AI we are (and we will) dealing with. An AI can indeed be classified in three ways: a narrow AI, which is nothing more than a specific domain application or task that gets better by ingesting further data and "learns" how to reduce the output error. An example here is Deep Blue for the chess game, but more generally this group includes all the functional technologies that serve a specific purpose. These systems are usually quite controllable because limited to specific tasks. When a program is instead not programmed for completing a specific task, but it could eventually learn from an application and apply the same bucket of knowledge to different environments, we face an Artificial General Intelligence (AGI). This is not *technology-as-a-service* as in the narrow case, but rather *technology-as-a-product*. The best example for this subgroup is Google DeepMind, although it is not a real AGI in all respects. We are indeed not there yet because even DeepMind cannot perform an intellectual task as a human would. In order to get there, much more progress on the brain structure functioning, brain processes optimization, and portable computing power development have to be made. Someone might think that an AGI can be easily achieved by piling up many narrow AIs, but in fact, this is not true: it is not a matter of number of specific skills a program can carry on, but rather the integration between all those abilities. This type of intelligence does not require an expert to work or to be tuned, as it would be the case for narrow AI, but it has a huge limitation: at the current state, it can be reached only through continuously streaming an infinite flow of data into the engine.

The final stage is instead called Super intelligent AI (ASI): this intelligence exceeds largely the human one, and it is able of scientific and creative thinking; it is characterized by general common wisdom; it has social skills and maybe an emotional intelligence. Although we often assume this intelligence to be a single super computer, it is more likely that it is going to be made by a network or a swarm of several intelligences.

The way in which we will reach the different stages is though still controversial, and many schools of thoughts exist. The symbolic approach claims that all the knowledge is symbolic and the representation space is limited, so everything should be stated in formal mathematical language. This approach has historically analyzed the complexity of the real world, and it had suffered at the same time from computational problems as well as understanding the origination of the knowledge itself. The statistical AI instead focuses on managing the uncertainty in the real world (Domingos et al. 2006), which lives in the inference realm contrarily to the more deductive logical AI. On a side then, it is not clear yet to what degree the human brain should be taken as an example: biological neural network seems to provide a great infrastructure for developing an AI, especially regarding the use of sparse distributed representations (SDRs) to process information.

3.2 A Bit of History

In spite of all the current hype, AI is not a new field of study, but it has its ground in the fifties. If we exclude the pure philosophical reasoning path that goes from the Ancient Greek to Hobbes, Leibniz, and Pascal, AI as we know it has been officially founded in 1956 at Dartmouth College, where the most eminent experts gathered to brainstorm on intelligence simulation. This happened only a few years after Asimov set his own three laws of robotics, but more relevantly after the famous paper published by Turing (1950), where he proposes for the first time the idea of a thinking machine and the more popular Turing test to assess whether such machine shows, in fact, any intelligence. As soon as the research group at Dartmouth publicly released the contents and ideas arisen from that summer meeting, a flow of government funding was reserved for the study of creating an intelligence that was not human.

At that time, AI seemed to be easily reachable, but it turned out that was not the case. At the end of the sixties, researchers realized that AI was indeed a tough field to manage, and the initial spark that brought the funding started dissipating. This phenomenon, which characterized AI along its all history, is commonly known as "AI effect", and is made of two parts: first, the constant promise of a real AI coming in the next ten years; and second, the discounting of behavior of AI after it mastered a certain problem, redefining continuously what intelligent means.

In the United States, the reason for DARPA to fund AI research was mainly due to the idea of creating a perfect machine translator, but two consecutive events wrecked that proposal, beginning what it is going to be called later on the first AI winter. In fact, the Automatic Language Processing Advisory Committee (ALPAC) report in US in 1966, followed by the "Lighthill report" (1973), assessed the feasibility of AI given the current developments and concluded negatively about the possibility of creating a machine that could learn or be considered intelligent. These two reports, jointly with the limited data available to feed the algorithms, as well as the scarce computational power of the engines of that period, made the field collapsing and AI fell into disgrace for the entire decade.

In the eighties, though, a new wave of funding in UK and Japan was motivated by the introduction of "expert systems", which basically were examples of narrow AI as above defined. These programs were, in fact, able to simulate skills of human experts in specific domains, but this was enough to stimulate the new funding trend. The most active player during those years was the Japanese government, and its rush to create the fifth-generation computer indirectly forced US and UK to reinstate the funding for research on AI.

This golden age did not last long, though, and when the funding goals were not met, a new crisis began. In 1987, personal computers became more powerful than Lisp Machine, which was the product of years of research in AI. This ratified the start of the second AI winter, with the DARPA clearly taking a position against AI and further funding.

Luckily enough, in 1993 this period ended with the MIT Cog project to build a humanoid robot, and with the Dynamic Analysis and Replanning Tool (DART)—that paid back the US government of the entire funding since 1950—and when in 1997 DeepBlue defeated Kasparov at chess, it was clear that AI was back to the top.

In the last two decades, much has been done in academic research, but AI has been only recently recognized as a paradigm shift. There are of course a series of causes that might bring us to understand why we are investing so much into AI nowadays, but there is a specific event we think it is responsible for the last five-years trend. If we look at Fig. 3.1, we notice that regardless all the developments achieved, AI was not widely recognized until the end of 2012. The figure has been indeed created using CB Insights Trends, which basically plots the trends for specific words or themes (in this case, Artificial Intelligence and Machine Learning).

More in details, I draw a line on a specific date I thought to be the real trigger of this new AI optimistic wave, i.e., Dec. 4th 2012. That Tuesday, a group of researchers presented at the Neural Information Processing Systems (NIPS) conference detailed information about their convolutional neural networks that granted them the first place in the ImageNet Classification competition few weeks before (Krizhevsky et al. 2012). Their work improved the classification algorithm from 72 to 85% and set the use of neural networks as fundamental for artificial intelligence. In less than two years, advancements in the field brought classification in the ImageNet contest to reach an accuracy of 96%, slightly higher than the human one (about 95%). The picture shows also three major growth trends in AI development, outlined by three major events: the 3-years-old DeepMind being acquired by Google in Jan. 2014; the open letter of the Future of Life Institute signed by more than 8000 people and the study on reinforcement learning released by Deepmind (Mnih et al. 2015) in February 2015; and finally, the paper published on Nature in Jan. 2016 by DeepMind scientists on neural networks (Silver et al. 2016) followed by the impressive victory of AlphaGo over Lee Sedol in March.

Fig. 3.1 Artificial intelligence trend for the period 2012–2016

AI is intrinsically highly dependent on funding because it is a long-term research field that requires an immeasurable amount of effort and resources to be fully depleted. There are then raising concerns that we might currently live the next peak phase (Dhar 2016), but also that the thrill is destined to stop soon. However, I believe that this new era is different for three main reasons: (i) (big) data, because we finally have the bulk of data needed to feed the algorithms; (ii) the technological progress, because the storage ability, computational power, better and greater bandwidth, and lower technology costs allowed us to actually make the model digesting the information they needed; and (iii) the resources democratization and efficient allocation introduced by Uber and AirBnb business models, which is reflected in cloud services (i.e., Amazon Web Services) and parallel computing operated by GPUs.

3.3 Why AI Is Relevant Today

The reason why we are studying AI right now more actively is clearly because of the potential applications it might have, because of the media and general public attention it received, as well as because of the incredible amount of funding investors are devoting to it as never before.

Machine learning is being quickly commoditized, and this encourages a more profound democratization of intelligence, although this is true only for low-order knowledge. If from one hand a large bucket of services and tools are now available to final users, on the other hand, the real power is concentrating into the hands of few major incumbents with the data availability and computational resources to really exploit AI to a higher level.

Apart from this technological polarization, the main problem the sector is experiencing can be divided into two key branches: first, the misalignments of (i) the long term AGI research sacrificed for the short-term business applications, and (ii) what AI can actually do against what people think or assume it does. Both the issues stem from the high technical knowledge intrinsically required to understand it, but they are creating hype around AI. Part of the hype is clearly justified, because AI has been useful in those processes that are historically hard to be automated because of the requirement of some degree of domain expertise.

Secondly, the tight relationship machine and humans have, and how they interact with each other. We are participating to an enormous cultural shift in the last few years because the human being was originally the creature in charge of acting, while the machine was the security device for unwanted scenarios. However, nowadays the roles have been inverted, and machines are often in charge while the humans are simply monitoring. Even more important, this relationship is changing our own being: people normally believe that machines are making humans more similar to them as humans are trying to do the same with computers, but there are thinkers who judge this cross-pollination as a way for humans to become even more humans (Floridi 2014). The only thing that seems to be commonly accepted is that fact that,

in order to shorten the AI adoption cycle, we should learn how to not trust our intuition all the time, and let the machine changing us either in a more human or more mechanical way.

So the natural question everyone is asking is "where machines stand with respect to humans?" Well, the reality is that we are still far from the point in which a superintelligence will exceed human intelligence—the so-called Singularity (Vinge 1993). The famous futurist Raymond Kurzweil proposed in 1999 the idea of the law of accelerating returns, which envisages an exponential technological rate of change due to falling costs of chips and their increasing computational capacity. In his view, the human progress is S-shaped with inflection points corresponding to the most relevant technological advancements, and thus proceeds by jumps instead of being a smooth and uniform progress. Kurzweil also borrowed Moore's law to estimate accurately the precise year of the singularity: our brain is able of 1016 calculations per second (cps) and 1013 bits of memory, and assuming Moore's law to hold, Kurzweil computed we will reach an AGI with those capabilities in 2030, and the singularity in 2045.

I believe though this is a quite optimistic view because the intelligence the machines are provided with nowadays is still only partial. They do not possess any common sense, they do not have any sense of what an object is, they do not have any earlier memory of failed attempts, they are not conscious—the so-called the "Chinese room" argument, i.e., even if a machine can perfectly translate Chinese to English and vice versa, it does not really understand the content of the conversation. On the other side, they solve problems through structured thinking, they have more storage and reliable memory, and raw computational power. Humans instead tried to be more efficient and select ex-ante data that could be relevant (at the risk of losing some important information), they are creative and innovative, and extrapolate essential information better and faster from only a few instances, and they can transfer and apply that knowledge to unknown cases. Humans are better generalists and work better in an unsupervised learning environment. There are easy intuitive tasks almost impossible for computer (what humans do "without thinking"), while number-intensive activities are spectacularly easy for a machine (the "hard-thinking" moments for our brain)—in other words, activities essential for survival that have to be performed without effort are easier for human rather than for machines. Part of this has been summarized by Moravec's paradox with a powerful statement: high-level reasoning requires little computation, and it is then feasible for a machine as well, while very simple low-level sensorimotor skills would demand a gigantic computational effort.

All the considerations made so far do not end in themselves but are useful to sketch the important design aspects to be taken into account when building an AI engine. In addition to those, few characteristics emerged as fundamental for progressing toward an AGI: robustness, safety, and hybridization. As intended in Russell et al. (2015), an AI has to be verified (acting under formal constraints and conforming to formal specifications); validated (do not pursue unwanted behaviors under the previous constraints); secure (preventing intentional manipulation by third parties, either outside or inside); and controlled (humans should have ways to

reestablish control if needed). Second, it should be safe according to Igor Markov's view: AI should indeed have key weaknesses; self-replication of software and hardware should be limited; self-repair and self-improvement should be limited; and finally, access to energy should be limited. Last, an AI should be created through a hybrid intelligence paradigm, and this might be implemented following two different paths: letting the computer do the work, and then either calling in humans in for ambiguous situations or calling them to make the final call. The main difference is that the first case would speed things up putting the machines in charge of deciding (and would use humans as feedback) but it requires high data accuracy.

The conclusion of this first section can then be summarized as follows: AI is coming, although not as soon as predicted. This AI spring seems to be different from previous phases of the cycle for a series of reasons, and we should dedicate resources and effort in order to build an AI that would drive us into an optimistic scenario.

References

Bostrom, N. (2014). *Superintelligence: Paths, dangers, strategies*. Oxford: OUP Oxford.

Brennan, T. J., & Lo, A. W. (2011). The origin of behavior. *Quarterly Journal of Finance*, *7*, 1043–1050.

Brennan, T. J., & Lo, A. W. (2012). An evolutionary model of bounded rationality and intelligence. *PLoS ONE*, *7*(11), e50310.

Dhar, V., (2016). The future of artificial intelligence. In *Big data*, *4*(1), 5–9.

Domingos, P., Kok, S., Poon, H., Richardson, M., & Singla, P. (2006). Unifying logical and statistical AI. In *Proceeding of the 21st National Conference on Artificial Intelligence.* (Vol. 1, pp. 2–7).

Floridi, L. (2014). *The fourth revolution: How the infosphere is reshaping human reality*. Oxford: OUP Oxford.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems* (pp. 1097–1105).

Kurzweil, R. (1999). *The age of spiritual machines: When computers exceed human intelligence*. London: Penguin Books.

Lighthill, J. (1973). Artificial intelligence: A general survey. In *Artificial intelligence: A paper symposium*, Science Research Council.

Lo, A. W. (2012). Adaptive markets and the new world order. *Financial Analysts Journal*, *68*(2), 18–29.

Lo, A. W. (2013). The origin of bounded rationality and intelligence. In *Proceedings of the American Philosophical Society*, *157*(3), 269–280.

Mnih, V., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, *518*, 529–533.

Müller, V. C., & Bostrom, N. (2016). Future progress in artificial intelligence: A survey of expert opinion. In V. C. Müller (Ed.), *Fundamental issues of artificial intelligence* (pp. 553–571). Berlin: Springer.

Russell, S., Dewey, D., & Tegmark, M. (2015). Research priorities for robust and beneficial artificial intelligence. *AI Magazine, 36*(4), 105–114.

Silver, D., et al. (2016). Mastering the game of go with deep neural networks and tree search. *Nature, 529,* 484–489.

Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, *49*, 433–460.

Vinge, V. (1993). The coming technological singularity: How to survive in the post-human era. In NASA. Lewis Research Center, *Vision 21: Interdisciplinary science and engineering in the era of cyberspace* (pp. 11–22).

Chapter 4
AI Knowledge Map: How to Classify AI Technologies

We have seen AI is a broad field and encompasses multiple type of technologies. The ways in which they could be classified are very different, but in this chapter, we will try to provide a new visual form to capture all the most relevant AI-related technologies.

Working with strategic innovation agency Axilo, we wanted to create a **visual tool for people to grasp at a glance the complexity and depth** of this toolbox, as well as laying down a map that could help people orientating in the AI jungle. You should look at the following graph as a way to **organize unstructured knowledge into a sort of ontology** with the final aim not to accurately represent all the existing information on AI but rather to have a tool **to describe and access part of that information set**.

What follows in Fig. 4.1 is then an effort to draw an **architecture to access knowledge** on AI and follow emergent dynamics, a gateway on pre-existing knowledge on the topic that will allow you to scout around for additional information and eventually create new knowledge on AI.

The utility of the final work should therefore help you achieve three things: **making sense** of what is going on and have a map to follow the path; understanding **where machine intelligence is used today** (with respect to where was not used for in the past); understanding what and how many **problems are reframed** to make possible for AI to tackle them (if you are familiar with the work of Agrawal et al. 2018 those are direct consequences of the drop in cost of *prediction technologies*).

So let's jump to **the AI Knowledge Map (AIKM)** now.

On the axes, you will find two macro-groups, i.e., the AI Paradigms and the AI Problem Domains. The **AI Paradigms** (X-axis) are really the approaches used by

---

Fig. 4.1 AI knowledge map

AI researchers to solve specific AI-related problems (it does include the approaches we are aware of up to date). On the other side, the **AI Problem Domains** (Y-axis) are historically the type of problems AI can solve. In some sense, it also indicates the potential capabilities of an AI technology.

Hence, I have identified the following the AI paradigms:

– **Logic-based tools**: tools that are used for knowledge representation and problem-solving;
– **Knowledge-based tools**: tools based on ontologies and huge databases of notions, information, and rules;
– **Probabilistic methods**: tools that allow agents to act in incomplete information scenarios;
– **Machine learning**: tools that allow computers to learn from data;
– **Embodied intelligence**: engineering toolbox, which assumes that a body (or at least a partial set of functions such as movement, perception, interaction, and visualization) is required for higher intelligence;
– **Search and optimization**: tools that allow intelligently searching through many possible solutions.

Those six paradigms also fall into three different macro-approaches, namely *Symbolic*, *Sub-symbolic* and *Statistical* (represented by different colors). Briefly, the **Symbolic approach** states that human intelligence could be reduced to symbol

manipulation, the **Sub-symbolic** one that no specific representations of knowledge should be provided ex-ante, while the **Statistical approach** is based on mathematical tools to solve specific sub-problems.

A quick additional note: you might hear people talking about "**AI tribes**", a concept proposed by Pedro Domingos (2015) that clusters researchers in groups based on the approaches they use to solve problems. You can easily map those five tribes with our paradigm classification (not considering the embodied intelligence group), i.e. **Symbolists** with Logic-based approach (they use logical reasoning based on abstract symbols); **Connectionists** with Machine learning (they are inspired by the mammalian brain); **Evolutionaries** with Search and Optimization (they are inspired by the Darwinian evolution); **Bayesians** with Probabilistic methods (they use probabilistic modeling); and finally **Analogizers** with Knowledge-based methods, since they try to extrapolate from existing knowledge and previous similar cases.

The vertical axis instead lays down the problems AI has been used for, and the classification here is quite standard:

– **Reasoning**: the capability to solve problems;
– **Knowledge**: the ability to represent and understand the world;
– **Planning**: the capability of setting and achieving goals;
– **Communication**: the ability to understand language and communicate;
– **Perception**: the ability to transform raw sensorial inputs (e.g., images, sounds, etc.) into usable information.

I am still interrogating myself whether this classification is large enough to capture all the spectrum of problems we are currently facing or whether more instances should be added (e.g., *Creativity* or *Motion*). For the time being though, I will stick with the 5-clusters one.

The patterns of the boxes instead divide the technologies into two groups, i.e., **narrow applications** and **general applications**. The words used are on purpose slightly misleading but bear with me for one second and I will explain what I meant. For anyone getting started in AI, knowing the difference between **Weak/Narrow AI** (*ANI*), **Strong/General AI** (*I*), and **Artificial Super Intelligence** (*ASI*) is paramount. For the sake of clarity, ASI is simply a speculation up to date, General AI is the final goal and holy grail of researchers, while narrow AI is what we really have today, i.e., a set of technologies which are unable to cope with anything outside their scope (which is the main difference with AGI).

The two types of lines used in the graph (continuous and dotted) then want to explicitly point to that distinction and make you confident that when you will read some other introductory AI material you won't be completely lost. However, at the same time, the difference here outlines technologies that can only *solve a specific task* (usually better than humans—**Narrow** applications) and others that *can today or in the future solve multiple tasks* and interact with the world (better than many humans—**General** applications).

Finally, let's see what there is within the graph itself. In the map, the different classes of AI technologies are represented. Note, I am intentionally not naming specific algorithms but rather clustering them into macro-groups. I am not either providing you with a value assessment of what it works and what it does not, but simply listing what researchers and data scientists can tap into.

So how do you read and interpret the map? Well, let me give you two examples to help you do that. If you look at Natural Language Processing, this embeds a class of algorithms that use a combination of a knowledge-based approach, machine learning and probabilistic methods to solve problems in the domain of perception. At the same time though, if you look at the blank space at the intersection between Logic-based paradigm and Reasoning problems, you might wonder why there are not technologies there. What the map is conveying is not that it does not categorically exist a method that can fill up that space, but rather that when people approach a reasoning problem they rather prefer to use a Machine Learning approach, for instance.

To conclude this explanation, this is the full list of technologies included with their own definitions:

– **Robotic Process Automation (RPA)**: technology that extracts the list of rules and actions to perform by watching the user doing a certain task;
– **Expert Systems**: a computer program that has hard-coded rules to emulate the human decision-making process. **Fuzzy systems** are a specific example of rule-based systems that map variables into a continuum of values between 0 and 1, contrary to traditional digital logic which results in a 0/1 outcome;
– **Computer Vision (CV)**: methods to acquire and make sense of digital images (usually divided into *activities recognition*, *images recognition*, and *machine vision*);
– **Natural Language Processing (NLP)**: sub-field that handles natural language data (three main blocks belong to this field, i.e., *language understanding*, *language generation*, and *machine translation*);
– **Neural Networks (NNs or ANNs)**: a class of algorithms loosely modeled after the neuronal structure of the human/animal brain that improves its performance without being explicitly instructed on how to do so. The two majors and well-known sub-classes of NNs are **Deep Learning** (a neural net with multiple layers) and **Generative Adversarial Networks** (GANs—two networks that train each other);
– **Autonomous Systems**: sub-field that lies at the intersection between robotics and intelligent systems (e.g., intelligent perception, dexterous object manipulation, plan-based robot control, etc.);
– **Distributed Artificial Intelligence (DAI)**: a class of technologies that solve problems by distributing them to autonomous "agents" that interact with each other. **Multi-agent systems (MAS)**, **Agent-based modeling (ABM)**, and **Swarm Intelligence** are three useful specifications of this subset, where collective behaviors emerge from the interaction of decentralized self-organized agents;

– **Affective Computing**: a sub-field that deal with emotions recognition, interpretation, and simulation;
– **Evolutionary Algorithms (EA)**: it is a subset of a broader computer science domain called evolutionary computation that uses mechanisms inspired by biology (e.g., mutation, reproduction, etc.) to look for optimal solutions. **Genetic algorithms** are the most used sub-group of EAs, which are search heuristics that follow the natural selection process to choose the "fittest" candidate solution;
– **Inductive Logic Programming (ILP)**: sub-field that uses formal logic to represent a database of facts and formulate hypothesis deriving from those data;
– **Decision Networks**: is a generalization of the most well-known **Bayesian networks**/inference, which represent a set of variables and their probabilistic relationships through a map (also called *directed acyclic graph*);
– **Probabilistic Programming**: a framework that does not force you to hardcode specific variable but rather works with probabilistic models. **Bayesian Program Synthesis (BPS)** is somehow a form of probabilistic programming, where Bayesian programs write new Bayesian programs (instead of humans do it, as in the broader probabilistic programming approach);
– **Ambient Intelligence (AmI)**: a framework that demands physical devices into digital environments to sense, perceive, and respond with context awareness to an external stimulus (usually triggered by a human action).

In order to solve a specific problem, you might follow one or more approaches, that in turn means one or more technologies given that many of them are not at all mutually exclusive but rather complementary.

Finally, there is another relevant classification that I have not embedded into the graph above (i.e., the different type of analytics) but that is worth to be mentioned for the sake of completeness. You may actually encounter five distinct types of analytics: *descriptive analytics* (what happened); *diagnostic analytics* (why something happened); *predictive analytics* (what is going to happen); *prescriptive analytics* (recommending actions); and *automated analytics* (taking actions automatically). You might also be tempted to use it to somehow classify the technologies above, but the reality is that this is a functional classification and a process one rather than a product one—in other words, every technology in the spectrum can fulfill those five analytics functions.

References

Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines*: The simple economics of artificial intelligence. Boston: Harvard Business Review Press.
Domingos, P. (2015). *The master algorithm: How the quest for the ultimate learning machine will remake our world*. New York: Basic Books.

Chapter 5
Advancements in the Field

AI is moving at a stellar speed and is probably one of most complex and present sciences. The complexity here is not meant as a level of difficulty in understanding and innovating (although of course, this is quite high), but as the degree of inter-relation with other fields apparently disconnected.

There are basically two schools of thought on how an AI should be properly built: the Connectionists start from the assumption that we should draw inspiration from the neural networks of the human brain, while the Symbolists prefer to move from banks of knowledge and fixed rules on how the world works. Given these two pillars, they think it is possible to build a system capable of reasoning and interpreting.

In addition, a strong dichotomy is naturally taking shape in terms of problem-solving strategy: you can solve a problem through a simpler algorithm, which though it increases its accuracy in time (iteration approach), or you can divide the problem into smaller and smaller blocks (parallel sequential decomposition approach).

Up to date, there is not a clear answer on what approach or school of thoughts works the best, and thus I find appropriate to briefly discuss major advancements in both pure machine learning techniques and neuroscience with an agnostic lens.

5.1 Machine Learning

Machine learning techniques can be roughly divided into supervised methods and unsupervised methods, with the main difference of whether the data are labelled (supervised learning) or not (unsupervised). A third class can be introduced when we talk about AI: reinforcement learning (RL). RL is a learning method for machines based on the simple idea of reward feedback: the machine indeed acts in a specific set of circumstances with the goal of maximizing the potential future (cumulative) reward. In other words, it is a trial-and-error intermediate method

between supervised and unsupervised learning: the data labels are indeed assigned only after the action, and not for every training example (i.e., they are sparse and time-delayed). RL usually comes with two major problems, namely the credit assignment problem and the explore-exploit dilemma—plus a series of technical issues such as the curse of dimensionality, non-stationary environments, or partial observability of the problem. The former one concerns the fact that rewards are, by definition, delayed, and you might need a series of specific actions in order to achieve your goal. The problem is then to identify which of the preceding action was actually responsible for the final output (and to get the reward then), and if so to what degree. The latter problem is instead an optimal searching problem: the software has to map the environment as accurately as possible in order to figure out its reward structure. There is an optimal stop problem—a sort of satisficing indeed: to what extent the agent should keep exploring the space to look for better strategies, or start exploiting the ones it already knows (and knows that work)?

In addition to the present classification, machine learning algorithms can be classified based on the output they produce: classification algorithms; regressions; clustering methods; density estimation; and dimensionality reduction methods.

The new AI wave encouraged the development of innovative ground-breaking techniques, as well as it brought back to the top a quite old concept, i.e., the use of artificial neural networks (ANNs).

Artificial Neural Networks are a biologically-inspired approach that allows software to learn from observational data—in this sense sometimes is said they mimic the human brain. The first ANN named Threshold Logic Unit (TLU) was introduced in the Forties by McCulloch and Pitts (1943), but only forty years later Rumelhart et al. (1986) pushed the field forward designing the back-propagation training algorithm for feed-forward multi-layer perceptrons (MLPs).

The standard architecture for any ANNS is having a series of nodes arranged in an input layer, an output layer, and a variable number of hidden layers (that characterize the depth of the network). The inputs from each layer are multiplied by a certain connection weight and summed up, to be compared to a threshold level. The signal obtained through the summation is passed into a transfer function, to produce an output signal that is, in turn, passed as input into the following layer. The learning happens in fact in the multiple iterations of this process, and it is quantitatively computed by choosing the weighting factors that minimize the input-output mapping error given a certain training dataset.

ANNs do not require any prior knowledge to be implemented, but on the other side, they can still be fooled because of it. They are often also called Deep Learning (DL), especially for the case in which there are many layers that perform computational tasks. There exist many types of ANNs up to date, but the most known ones are Recurrent Neural Networks (RNNs); Convolutional Neural Networks (CNNs); and Biological Neural Networks (BNNs).

RNNs use the sequential information to make accurate predictions. In traditional ANNs, all the inputs are independent one from the other. RNNs perform instead a certain task for every element of the sequence, keeping a sort of memory of the previous computations. CNNs try instead to mirror the structure of the mammalian

visual cortex and they have every layer working as detection filters for detecting specific patterns in the original data (and this is why they are really suitable for object recognition). Finally, BNNs are more a sub-field of ANNs rather than a specific application. The best example of this class is, in my opinion, the Hierarchical Temporal Memory (HTM) model developed by Hawkins and George of Numenta, Inc., which is a technology that captures both the structural and algorithmic properties of the neocortex.

In spite of the big hype around deep learning possibilities, all that glitters is not gold. DL is for sure a great step ahead toward the creation of an AGI, but it also presents limitations. The greatest one is the exceptional amount of data required to work properly, which represents the major barrier to a wider cross-sectional application. DL is also not easy to debug, and usually, problems are solved by feeding more and more data into the network, which creates a tighter big-data-dependency. Furthermore, DL is quite useful to bring to light hidden connections and correlations but is not informative at all regarding the causation (the why of things).

The data need imposes a considerable amount of time to train a network. In order to reduce this time, networks are often trained in parallel, either partitioning the model across different machines on different GPU cards (model parallelism) or reading different (random) buckets of data through the same model run on different machines to tune the parameters (data parallelism).

Because of the limitations just mentioned, a series of other tools have been developed over the years. **Particle Swarm Optimization** (**PSO**) is a computational method that iteratively improves candidate solution to optimize a certain problem (Kennedy and Eberhart 1995). The initial population of candidates (namely dubbed particles) is moved around in the search-space, and it has single particles that optimize their own position both locally and with respect to the entire search-space—creating then an optimized swarm. **Agent-based Computational Economics** (**ACE**) is an additional tool that lets agents interacting according to pre-specified rules into simulated environments (Arthur 1994). Starting from some initial condition imposed by the modeler, the dynamic systems evolves over time as interactions between agents occur (and as they learn from previous interactions).

Evolutionary Algorithms (**EA**) are instead a broad class of techniques that find solutions to optimization problems through concepts borrowed from natural evolution, i.e., selection, mutations, inheritance, and crossover. An example of EA is the **Genetic Algorithm** (**GA**), which is an adaptive search heuristic that attempts to mimic the natural selection process (Holland 1975). It is an evolutionary computing search optimization method that starts from a base population of candidate solutions and makes them evolving according to the "survival of the fittest" principle. **Genetic Programming** (**GP**) is an extension of GA (Koza 1992) because it basically applies a GA to a population of computer programs. It creates the chromosomes (i.e., the initial population of programs) made by a predefined set of functions and a set of terminals, and it randomly combines them into a tree-structure. In this context, the previous terminology acquires a slightly different connotation: reproduction means copying another computer model from existing

population; cross-over means randomly recombining chosen parts of two computer programs, and mutation is a random replacement of chosen functional or terminal node. **Evolutionary Polynomial Regressions** (**EPRs**) are instead hybrid regressions that use GA to select the exponents of the polynomial, and a numerical regression (i.e., least square regression) to compute the actual coefficients (Giustolisi and Savic 2006). A final interesting model is called **Evolutionary Intelligence** (**EI**) or **Evolutionary Computation** (**EC**), and it has been recently developed by Sentient Technologies, LLC. It begins randomly generating trillions of candidate solutions (called genes) that by definition would probably perform poorly. They are then tested against training data, and a fitness score allowed the software to rank the best solutions (and eliminates the worst). Parts of the emerging candidates are then used to reassemble new populations, and the process restarts until a convergence is achieved.

To conclude this section, two additional approaches are worthy to be acknowledged. First, **Generative Models** (**GMs**) have been initially proposed by Shannon (1948), but recently brought back to the top by OpenAI, a non-profit AI research institute based in San Francisco (Salimans et al. 2016; Chen et al. 2016). This class of models is intuitively defined as those models we can randomly generate data for, assumed some hidden parameters. Once the data are feed, the system specifies a joint probability distribution and label sequences of data.

Second, Cao and Yang (2015) proposed a new method that converts the learning algorithm into a summation form, instead of proceeding directly from each training data point. It is called **Machine Unlearning** (**MU**), and it allows the systems to "forget" unwanted data. They actually introduce an intermediate layer of summation between the algorithm and the training data points, such that they will not depend on each other anymore, but only on the summations themselves. In this way, they learning process is much faster, and it can be updated incrementally without training again the model from scratch—which is quite time-intensive and costly. Hence, if some data and its lineage want to be eliminated, the system does not need to recompute the entire lineage anymore—a term coined by the two authors to indicate the entire data propagation network—but it can simply recompute a small number of summations.

5.2 Neuroscience Advancements

Along with the advancements in pure machine learning research, we have done many steps ahead toward a greater comprehension of the brain mechanisms. Although much has still to be understood, we have nowadays a slightly better overview of the brain processes, and this might help to foster the development of an AGI. It seems clear that try to fully mimic the human brain is not a feasible approach, and is not even the correct one. However, drawing inspiration from how the brain works is a completely different story, and the study of neuroscience could

both stimulate the creation of new algorithms and architectures, as well as validate the use of current machine learning research toward a formation of an AGI.

More in detail, according to Numenta's researchers AI should be inspired to the human neocortex. Although a common theoretical cortical framework has not been fully accepted by the scientific community, according to Numenta a cortical theory should be able to explain: (i) how layers of neurons can learn sequences; (ii) the properties of SDRs; (iii) unsupervised learning mechanism with streaming temporal data flows; (iv) layer to layer connectivity; (v) how brain regions model the world and create behaviors; and finally, (vi) the hierarchy between different regions. These can be seen then as the six principles any biological or artificial intelligence should possess to be defined as such. Intuitively, it sounds a reasonable model, because the neocortex learns from sensory data, and thus it creates a sensory-motor model of the world. Unfortunately, we do not fully comprehend how the neocortex works yet, and this demands a machine intelligence be created flexible as much as robust at the same time.

In a more recent work, Hawkins and Ahmad (2016) turned their attention on a neuroscientific problem who is though crucial to the development of an AGI. They tried to explain how neurons integrate inputs from thousands of synapses, and their consequent large-scale network behavior. Since it is not clear why neurons have active dendrites, almost every ANNs created so far do not use artificial dendrites at all, and this would suggest that something is probably missing in our artificial structures. Their theory explains how networks of neurons work together, assumed all the many thousands of synapses presented in our brain. Given those excitatory neurons, they proposed a model for sequence memory that is a universal characteristic of the neocortical tissue, and that if correct would have a drastic impact on the way we design and implement artificial minds.

Rocki (2016) also highlighted few aspects specifically relevant for building a biologically inspired AI—specifically, the necessary components for creating a general-purpose learning algorithm. It is commonly assumed that humans do not learn in a supervised way, but they learn (unsupervised) to interpret the input from the environment, and they filter out as much data as possible without losing relevant information (Schmidhuber 2015). Somehow, the human brain applies a sort of Pareto's rule (or a Minimum Description Length rule otherwise) to information it gathers through sensory representations, and keeps and stores only the information that can explain the most of what is happening. According to Rocki, unsupervised learning regularizes and compresses information making our brain a data compactor (Bengio et al. 2012; Hinton and Sejnowski 1999).

In addition to being unsupervised, Rocki hypotheses that the architecture of a general-learning algorithm has to be compositional; sparse and distributed; objectiveless; and scalable. Human brain learns sequentially, starting from simpler patterns and breaking up more complex problems in terms of those simpler bricks it already understood—and this type of hierarchy and compositional learning is indeed well captured by deep learning. As already pointed out by Ahmad and Hawkins (2015), sparse distributed representations are essential, and they are much more noisy-resistant than their dense counterparts. However, there are much more

peculiarities that make SDRs preferable: there are no region-specific algorithms in the brain, but the cortical columns act as independent feature detectors. Each column becomes active in response to a certain stimulus, and at the same time, it laterally inhibits other adjacent columns, forming thus sparse activity patterns. Since they are sparse, it is easier to reverse engineer a certain external signal and extract information from it (Candès et al. 2006). The property of being distributed helps instead in understanding the causes of patterns variations. SDRs also facilitates the process described above of filtering out useless information. They represent minimum entropy-codes (Barlow et al. 1989) that provide a generalized learning mechanism with simpler temporal dependencies.

The reason why the learning process should not have a clear stated objective is slightly controversial, but Rocki—and Stanley and Lehman (2015) before him—support this argument as the only way to achieve and form transferrable concepts. Moreover, Rocki states scalability as fundamental for a general-learning architecture. The brain is inherently a parallel machine, and every region has both computational and storing tasks (and this is why GPUs are much more efficient than CPUs in deep learning). This would suggest an AI to have a hierarchical structure that separates local learning (parallel) from higher-order connections (synapses updates), as well as a memory that can itself compute, in order to reduce the energy cost of data transfers.

Rocki eventually concludes with some further functional rather than structural ingredients for the formation of an AI, namely: compression; prediction; understanding; sensorimotor; spatiotemporal invariance; context update; and pattern completion. We discussed the importance of compression and sensorimotor before, and we can think of AGI as a general-purpose compressor that forms stable representations of abstract concepts—although this point is controversial according to the *no free lunch theorem* (Wolpert and Macready 1997) that indirectly states that this algorithm cannot exist. We can also see prediction as of a weak form of spatiotemporal coherence of the world, and then we can argue learning to predict to be equivalent to understanding. Finally, we need to incorporate a continuous loop of bottom-up predictions and top-down contextualization to our learning process, and this contextual spatiotemporal concept would also allow for a disambiguation in the case of multiple (contrasting) predictions.

5.3 Hardware and Chips

As we explained before, the recent surge of AI and its rapidly becoming a dominant discipline are partially due to the exponential degree of technological progress we faced over the last few years. What it is interesting to point out though is that AI is deeply influencing and shaping the course of technology as well.

First of all, the Graphics Processing Units (GPUs) have been adapted from traditional graphical user interface applications to alternative parallel computing operations. NVIDIA is leading this flow and is pioneering the market with the

CUDA platform and the recent introduction of Tesla P100 platform (the first GPU designed for hyperscale data center applications). On top of P100, they also created the first full server appliance platform (named DGX-1), which will bring deep learning to an entirely new level. Very recently, they also released the Titan X, which is the biggest GPU ever built (3584 CUDA cores).

In general, the most impressive developments we observed are related to chips, especially Neuromorphic Processing Units (NPUs) ideated to emulate the human brain. Specific AI-chips have been created by major incumbents: IBM has released in 2016 the TrueNorth chip, which it is claimed to work very similarly to a mammalian brain. The chip is made of 5.4 billion transistors, and it is able to simulate up 1 million neurons and 256 million neural connections. It is equipped with 4000 cores that have 256 inputs lines (the axons) and as much output lines (neurons), which send signals only when electrical charges achieve a determined threshold.

This structure is quite similar to the Neurogrid developed by Stanford, although the academic counterpart is made of 16 different chips instead of the single one proposed by the software colossus.

Google, on the other hand, announced the introduction design of an application-specific integrated circuit (ASIC) thought and tuned specifically for neural networks—the so-called Tensor Processing Unit (TPU). The TPU optimizes the performance per watt specifically for machine learning problems, and it both powers RankBrain (i.e., Google Search) and DeepMind (i.e., AlphaGO).

Intel is working on similar chips as well, i.e., the Xeon Phi chip series, and the latest release has been named Knights Landing (KNL). KNL has up to 72 cores, and instead of being a GPU, it can be a primary CPU that reduces the need to offload machine learning to co-processors.

Even Qualcomm has invested enormous resources in the Snapdragon 820, and eventually into the deep learning SDK Snapdragon Neural Processing Engine and their Zeroth Machine Intelligence Platform.

The cost for all those chips is huge (on the order of billions for R&D, and hundred thousand dollars as selling cost), and they are not viable for retail consumers yet but only thought for enterprise applications. The main exception to this major trend is the mass-scale commercial AI chip called Eyeriss, released earlier in 2016 by a group of researchers at MIT. This chip—made of 168 processing engines—has been built on a smartphone's power budget and thus is particularly energy-friendly, but it presents anyway computational limitations.

Even though this is a cost-intensive game, several startups and smaller companies are considerably contributing to the space: Numenta open-source NuPIC, a platform for intelligent computing, to analyze streaming data. Knowm, Inc. has brought memristor chips to the market, which is a device that can change its internal resistance based on electrical signals fed into it (and used as a non-volatile memory). KnuEdge (and its subsidiaries KnuPath) created LambaFabric, which runs on a completely innovative architecture different not only from traditional GPUs but also from TPUs. Nervana Systems released an ASIC with a new high-capacity and high-speed memory technology called High Bandwidth Memory. Horizon Robotics

is another company actively working in the space, as well as krtkl, which has produced a new low-cost dual-core ARM processor (FPGA, Wi-Fi, Bluetooth) named Snickerdoodle.

A final note has to be made in favor of Movidius, which introduced a completely new concept, i.e., an all-in-one USB for deep learning. Codenamed Fathom Neural Compute Stick, it contains a chip called Myriad 2, which has been thought in partnership with Google specifically to tackle down any advanced image recognition issue (but it has been used also to power drones and robots of a diverse kind).

References

Ahmad, S., & Hawkins, J. (2015). *Properties of sparse distributed representations and their application to hierarchical temporal memory*. arXiv:1503.07469.

Arthur, B. W. (1994). Inductive reasoning and bounded rationality. *American Economic Review, 84*(2), 406–411.

Barlow, H. B., Kaushal, T. P., & Mitchison, G. J. (1989). Finding minimum entropy codes. *Neural Computation, 1*(3), 412–423.

Bengio, Y., Courville, A. C., & Vincent, P. (2012). *Unsupervised feature learning and deep learning: A review and new perspectives*. CoRR. arXiv:abs/1206.5538.

Candès, E. J., Romberg, J. K., & Tao, T. (2006). Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences, 59*(8), 1207–1223.

Cao, Y., & Yang, J. (2015). Towards making systems forget with machine unlearning. *IEEE Symposium on Security and Privacy, 2015,* 463–480.

Chen, X., Duan, X., Houthooft, R., Schulman, J., Sutskever, I., & Abbeel, P. (2016). *InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets*. arXiv:1606.03657.

Giustolisi, O., & Savic, D. A. (2006). A symbolic data-driven technique based on evolutionary polynomial regression. *Journal of Hydroinformatics, 8*(3), 207–222.

Hawkins, J., & Ahmad, S. (2016). Why neurons have thousands of synapses, a theory of sequence memory in neocortex. *Frontiers in Neural Circuits, 10*, 23.

Hinton, G., & Sejnowski, T. (1999). *Unsupervised learning: Foundations of neural computation*. Cambridge: MIT Press.

Holland, J. H. (1975). *Adaptation in natural and artificial systems*. Cambridge: MIT Press.

Kennedy, J., & Eberhart, R. (1995). Particle swarm optimization. In *Proceedings of IEEE International Conference on Neural Networks* (pp. 1942–1948).

Koza, J. R. (1992). *Genetic programming: On the programming of computers by means of natural selection*. Cambridge: MIT Press.

McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics, 5,* 115–133.

Rocki, K. (2016). *Towards machine intelligence* (pp. 1–15). CoRR. arXiv:abs/1603.08262.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature, 323*, 533–536.

Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., & Chen, X. (2016). *Improved techniques for training GANs*. arXiv:1606.03498.

Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks, 61*, 85–117.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal, 27*(379–423), 623–656.

Stanley, K. O., & Lehman, J. (2015). *Why greatness cannot be planned—The myth of the objective*. Berlin: Springer International Publishing.

Wolpert, D. H., & Macready, W. G. (1997). No free lunch theorems for optimization. *Transactions on Evolutionary Computation, 1*(1), 67–82.

Chapter 6
AI Business Models

AI is introducing radical innovation even in the way we think about business, and the aim of this section is indeed to categorize different AI companies and business models.

It is possible to look at the AI sector as really similar in terms of business models to the biopharma industry: expensive and long R&D; long investment cycle; low-probability enormous returns; concentration of funding toward specific phases of development. There are anyway two differences between those two fields: the experimentation phase, that is much faster and painless for AI, and the (absent) patenting period, which forces AI to continuously evolve and to use alternative revenue models (e.g., freemium model).

If we look from the incumbents' side, we might notice two different nuances in their business models evolution. First, the growth model is changing. Instead of competing with emerging startups, the biggest incumbents are pursuing an aggressive acquisition policy. I named this new expansion strategy the "***DeepMind strategy***" because it has become extremely common after the acquisition of DeepMind operated by Google. The companies are purchased when they are still early stage, in their first 1–3 years of life, where the focus is more on people and pure technological advancements rather than revenues (AI is the only sector in which the pure team value exceeds the business one). They maintain elements of their original brand and retain the entire existing team ("*acqui-hire*"). Companies maintain full independence, both physically speaking—often they keep in place their original head-quarters—as well as operationally. This independence is so vast to allow them to pursue acquisition strategies in turn (DeepMind bought Dark Blue Labs and Vision Factory in 2014). The parent company uses the subsidiary services and integrates rather than replaces the existing business (e.g., Google Brain and Deepmind).

It seems then that the acquisition costs are much lower than the opportunity cost of leaving around many brains, and it works better to (over) pay for a company today instead of being cutting out a few years later. In this sense, these acquisitions

are pure real option tools: they represent future possible revenues and future possible underlying layers where incumbents might end up building on top of.

The second nuance to point out is the emerging of the open source model in the AI sector, which is quite difficult to reconcile with the traditional SaaS model. Many of the cutting-edge technologies and algorithms are indeed provided for free and can be easily downloaded. So why incumbents are paying huge money and startups are working so hard to give all away for free? Well, there are a series of considerations to be made here. First, AI companies and departments are driven by scientists and academics, and their mindset encourages sharing and publicly presenting their findings. Second, open sourcing raises the bar of the current state of the art for potential competitors in the field: if it is publicly noted what you can build with TensorFlow, another company that wants to take over Google should publicly prove to provide at least what TensorFlow allows. It also fosters use cases that were not envisioned at all by the providing company and set up those tools as underlying technology everything should be built on top of which. Releasing for free software that do not require presence of high-tech hardware is also a great way for: (i) lowering the adoption barrier to entry, and get traction on products that would not have it otherwise; (ii) troubleshooting, because many heads are more efficient in finding and fixing bugs as well as looking at things from a different perspective; (iii) (crowd) validating, because often the mechanism, rationales, and implications might not be completely clear; (iv) shortening the product cycle, because from the moment a technical paper is published or a software release it takes weeks to have augmentations of that product; (v) to create a competitive advantage in data creation/collection, in attracting talents, and creating additive products based on that underlying technology; and (vi) more importantly, to create a *data network effect*, i.e., a situation in which more (final or intermediate) users create more data using the software, which in turn make the algorithms smarter, then the product better, and eventually attract more users.

There are the many reasons why this model is working, even though there are advocates who claim incumbents to not really be maximally open (Bostrom 2016) and to only release technology somehow old to them. My personal view is that companies are getting the best out of spreading their technologies around without paying any costs and any counter effect: they still have unique large datasets, platform, and huge investments capacity that would allow only them to scale up.

Regardless the real reasons behind this strategy, the effect of this business model on the AI development is controversial. According to Bostrom (2016), in the short term a higher openness could increase the diffusion of AI. Software and knowledge are non-rival goods, and this would enable more people to use, build on top on previous applications and technologies at a low marginal cost, and fix bugs. There would be strong brand implications for companies too.

In the long term, though, we might observe less incentive to invest in research and development, because of free riding. Hence, there should exist a way to earn monopoly rents from ideas individuals generate. On other side, what stands on the positive side is that open research is implemented to build absorptive capacity (i.e., is a mean of building skills and keeping up with state of art); it might bring to extra

profit from owning complementary assets whose value is increased by new technologies or ideas; and finally, it is going to be fostered by individuals who want to demonstrate their skills, build their reputation, and eventually increase their market value.

Although these notes on the effect of open research on AI advancements in short versus long term, it is not clear where this innovation will be promoted. We are looking at the transition from universities, where historically innovation and research lie, to the industry. This is not a new concept, but it is really emphasized in AI context. It has been created a vicious circle, in which universities lost faculty and researchers to the benefit of private companies because they can offer a combination of higher salary, more interesting problems, relevant large unique datasets, and virtually infinite resources. This does not allow universities to train the next generation of Ph.D. students that would be in charge of fostering the research one step ahead. The policy suggestion is then to fund pure research institutes (e.g., OpenAI) or even research-oriented companies (as for instance Numenta) to not lose the invaluable contribution that pure research has given to the field.

Most of the considerations made so far were either general or specific to big players, but we did not focus on different startup business models. An early stage company has to face a variety of challenges to succeed, and usually, they might be financial challenges, commercial problems, or operational issues. AI sector is very specific with respect to each of them: from a financial point of view, the main problem regards the absence of several specialized investors that could really increase the value of a company with more than mere money. The commercial issues concern instead the difficulties in identifying target customers and trying head around the open source model. The products are highly new and not always understood, and there might be more profitable ways to release them. Finally, the operational issues are slightly more cumbersome: as abovementioned, large dataset and consistent upfront investments are essential and might be detrimental to a shorter-term monetization strategy. A solution to the data problem may be found in the "data trap" strategy, that in venture capitalist Matt Turck's words consists of offering (often for free) products that can initialize a data network effect. In addition, the user experience and the design are becoming tangibly relevant for AI, and this creates friction in early stage companies with limited resources to be allocated between engineers, business, and design.

All those problems can create two major cross-sectional problems: the likely event to run out of money before hitting relevant milestones toward the next investment, as well as whether pursuing specific business applications to break even instead of focusing on product development.

In terms instead of classifying different companies operating in the space, there might be several different ways to think around machine intelligence startups (e.g., the classification proposed by Bloomberg Beta investor Shivon Zilis in 2015 is very accurate and useful for this purpose). I believe though that a too narrow framework might be counterproductive given the flexibility of the sector and the facility of transitioning from one group to another, and so I preferred to create a four-major-clusters categorization:

(i) **Academic spin-offs**: these are the more long-term research-oriented companies, which tackle problems hard to break. The teams are usually really experienced, and they are the real innovators who make breakthroughs that advance the field;

(ii) **Data-as-a-service (DaaS)**: in this group are included companies which collect specific huge datasets, or create new data sources connecting unrelated silos;

(iii) **Model-as-a-service (MaaS)**: this seems to be the most widespread class of companies, and it is made of those firms that are commoditizing their models as a stream of revenues. They can appear in three different forms:

  a. **Narrow AI**—a company that focus on solving a specific problem through new data, innovative algorithms, or better interfaces;
  b. **Value extractor**—a company that uses its models to extract value and insights from data. The solutions usually provided might either integrate with the clients' stack (through APIs or building specifically on top of customers' platform) or otherwise full-stacks solutions. All the models offered can be trained (operative models) or to be trained (raw models);
  c. **Enablers**—a company that is enabling the final user to do either her own analysis (all-in-one platforms), or allowing companies to make daily workflows more efficient, or eventually unlocking new opportunities through the creation of intermediate products (e.g., applications).

(iv) **Robot-as-a-service (RaaS)**: this class is made by virtual and physical agents that people can interact with. Virtual agents and chatbots cover the low-cost side of the group, while physical world systems (e.g., self-driving cars, sensors, etc.), drones, and actual robots are the capital and talent-intensive side of the coin.

The results of this categorization can be summarized into the following matrix, plotting the groups with respect to short-term monetization (STM) and business defensibility (Fig. 6.1).

Starting from the more viable products, the MaaS are the companies with the highest potential to monetize their products in the short term, but also the less defensible. DaaS on the other side is way less replicable, and highly profitable anyway. Academic spin-offs are the long bet, which is based on solid scientific research that makes them unique but not valuable form day one. Finally, RaaS companies are the ones who might face more problems, because of high obsolescence in hardware components and difficulties in creating the right interactive interfaces. This classification is not intended to rank any business based on how good they are, and it does not imply that specific companies belonging to specific classes are not going to be profitable or successful (e.g., X.ai is a high profitable company with a great product into the RaaS area). It is nothing more than a generalization tool useful to look at the sector through the correct lenses.
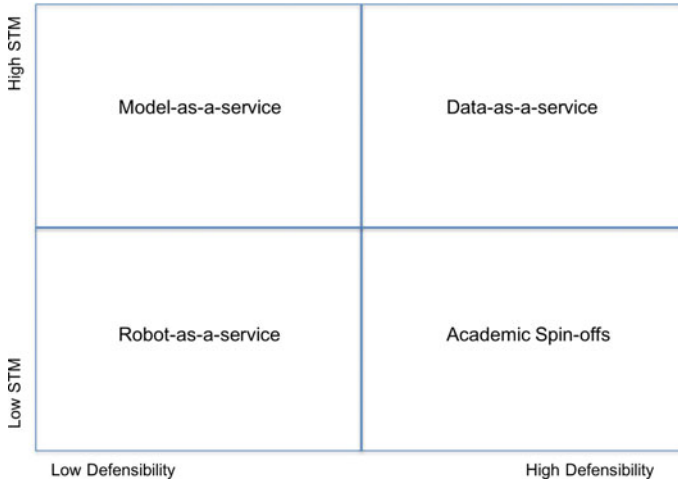
Fig. 6.1 Artificial intelligence classification matrix

To conclude this section, I want to highlight three final characteristics that *AI-as-a-technology* is introducing. First of all, AI is disrupting the traditional IoT business model because is bringing analytics to final customers instead of centralizing it.

Second, it is forcing business to keep customers in the loop—and the reasons why are manifold: (i) establishing trust into the product and the company; (ii) increasing the clients' retention building customary behaviors; (iii) improve sensibly the product through feedbacks.

The shifting focus on the final user as part of the product development is quickly becoming essential, to such a point that it represents a new business paradigm, i.e., the "***Paradigm 37–78***". I named this new pattern after the events of March 2016, in which AlphaGo defeated Lee Sedol in the Go game. In the move 37, AlphaGo surprised Lee Sedol with a move that no human would have ever tried or seen coming, and thus it won the second game. Lee Sedol rethought about that game, getting used to that kind of move and building the habit of thinking with a new perspective. He started realizing (and trusting) that the move made by the machine was indeed superb, and in game four he surprised in turn AlphaGo at Move 78 with something that the machine would not expect any human to do.

The Paradigm 37–78 is indeed a way to acknowledge that the users are the real value driver for building an effective AI engine: we make the machine better, and they make us better off in turn.

The last feature AI is changing is the way we think about data. First, AI is pushing business to question whether the information is always good and if the benefits linearly increase with a higher volume. This aspect is really important because AI is trained on data that have to be high quality to be effective (and this is why Twitter turned Microsoft's bot into a Hitler-loving sex robot). It is also forcing us to reflect on storing data that matter (rather than storing just for the sake of doing it), and to use

correctly data exhaust, i.e., those data generated as a by-product of online actions—in other words, they are not business core data, and they are by definition multiplicative with respect to the initial information (and thus much bigger). Finally, AI necessities are clearly underlining the cost-benefit trade-off of the inversely related relationship between accuracy and implementation time (either time to train the model, or time to produce the results and provide answers). The discussion on this specific topic is highly dependent on the sector and problem tackled: there are cases in which it is better the dollar cost of learning is largely overcome from higher accuracy, while others in which faster and responsive answers are way better than an incredibly accurate one.

Data is by far the perfect good: it does not deteriorate over time and can be reused; it is multipurpose; it multiplies by using or sharing. It is clearly up to date one of the greatest sources of competitive advantage for any machine learning firm, which represents also a problem: data polarization might result into few companies that channel and attract most of the data traffic, and other ones being (almost) completely excluded. In few years, this exponential trend might generate an enormous barrier to entry for the sector, compelling companies to create strategic partnerships with incumbents.

Fortunately, there are already stealth-mode companies working on reducing the dependency of AI on extremely large datasets (such as Vicarious or Geometric Intelligence for example): machines should indeed be able to learn from just a few instances as humans do. It is also not a coincidence that they are led by academics, because if the solution for the business is feeding the model with more data (narrowing down the bottleneck), for academics is instead focusing on transforming the algorithms for the better, and laying the foundation for the next evolutionary step.

Reference

Bostrom, N. (2016). *Strategic implications of openness in AI development*. Working paper.

Chapter 7
Hiring a Data Scientist

All this confusion and vagueness around definitions and concepts, and the hurdles technicalities of the big data black box have turned the people who analyze huge datasets into some kind of mythological figures. These people, who possess all the skills and the willingness to crunch numbers and providing insights based on them, are usually called data scientists. They have inherited their faith in numbers from the Pythagoreans before them, so it may be appropriate to name them *Datagoreans*. Their school of thinking, the Datagoreanism, encourages them to pursue the truth through data, and to exploit blending and fruitful interactions of different fields and approaches for postulating new theories and identifying hidden connections.

However, the general consensus about who they are and what they are supposed to do (and internally deliver) is quite loose. Just by browsing job offers for data scientists one understands that employers do not often really know what they are exactly looking for, and this is probably why everyone is complaining about the shortage of data scientists in the job market nowadays (Davenport and Patil 2012).

In reality, data scientists as imagined by most do not exist because it is a complete new figure, especially for the initial degrees of seniority. However, the proliferation of boot camps and structured university programs on one hand, and the companies increased awareness about this field on the other hand, will drive the job market towards its demand-supply equilibrium: firms will understand what they actually need in term of skills, and talents will be eventually able to provide those (verified) required abilities.

It is then necessary now to outline this new role, which is still half scientist half designer, and it includes a series of different skillsets and capabilities, akin to the mythological chimera. The pro ling is then provided in the following table, and it merges basically five different job roles into one as shown in Fig. 7.1: the computer scientist, the businessman, the statistician, the communicator, and the domain expert (a more complete list of skills could be found in the Appendix III).

Clearly, it is very cumbersome if not impossible to substitute five different people with a single one. This consideration allows us to draw several conclusions. First, collapsing five job functions into one is efficient on one hand because the
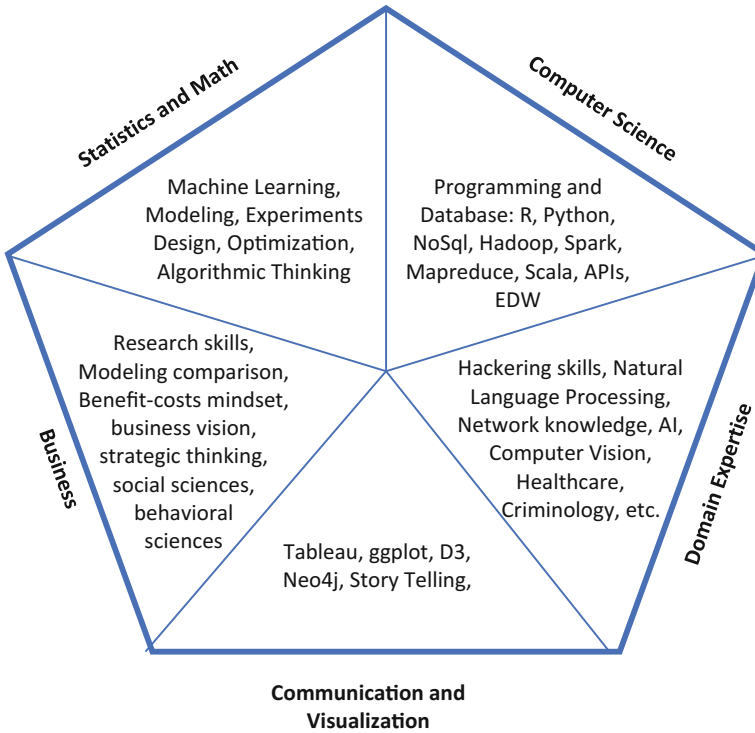
Fig. 7.1   Data scientist core skills set

value chain is concentrated and not dispersed, but on the other hand it can demand more time and resources. Indeed, a single individual can probably be less productive than five different people working on the same problem in the same timeframe. Secondly, hiring one specialist should cost less than a total of five semi-specialists, but much more than anyone of them singularly considered because of his specialization and high-level knowledge and flexibility (for the full salary toy-model, see Corea 2016). Data scientists seem to be fairly compensated in absolute terms, but their remuneration is definitely lower if compared to the cost structure they are facing to become such specialized figure. Indeed, the learning cost is really high because there are not so many designed programs for data science, so they have to do great efforts in filling their knowledge gaps. Furthermore, the universities and training monetary costs are really burdensome, and the opportunity costs quite heavy, and since the path is definitely new and not well established the choice of becoming a data scientist is risky and expensive.

All the considerations drawn so far point to a few suggestions for hiring data scientists: first of all, data science is teamwork, not a solo sport. It is important to hire different figures as part of the team, rather than exclusively for individual abilities.

Moreover, if a data science team is a company priority, the data scientists *have to be hired to stay*, because managing big data is a marathon rather than 100 m dash.

Secondly, data scientists come with two different DNAs: the scientific and the creative one. For this reason, they should be let free to learn and continuously study from one hand (the science side), and to create, experiment, and fail from the other (the creative side). They will never grow systematically and at a fixed pace, but they will do that organically based on their inclinations and multi-faceted nature. It is recommended to leave them with some spare time to follow their inspirations (some company is already doing that since years, and it consists in leaving them 10–20% of their working hours to pursue their personal ideas, to innovate, or simply for self-discovering). Furthermore, they have to be highly motivated, because often money even relevant is not anyway the crucial aspect to them. A high salary is indeed a signal of the respect the company has for their work—both past and future—and it is of course an incentive because they could potentially go working anywhere else. Although, the retention power of a good salary is quite low with respect to interesting daily challenges, and in order to align the data experts' interests with the company vision they have to be continuously fed with stimulating problems, and their work has to be relevant and impactful. Remember also that the scientist part requires them to be part of a bigger community, as well as the freedom to share concepts, ideas, and eventually working in parallel with peers. Even though the companies believe in patenting and scarce divulgation of what they do to maintain a sustainable competitive advantage, they have to compromise with the fact that scientists need to publish their research, sharing data, materials, and ideas.

Finally, do not be closed-minded and suppress any prejudice. Even if on percentage there is a good share of American male with a Ph.D. working in data science (King and Magoulas 2015), this may be indicative but not conclusive on the ideal candidate to hire: value the skills and capabilities more than titles or formal structured education—at least as soon as the field would be deeply-rooted and university degrees would be a good signal for skills owned. So far, in order to become a data scientist, the paths to be followed could be unconventional and various, so it is important to assess the abilities instead of deciding based on the type of background or degree level. Never forget that one of the real extra-value added by data science is different field contaminations and cross-sectional applications. It is also essential to take into account that not all the data specialists are the same (Liberatore and Luo 2012; Kandel et al. 2012), and it is possible to cluster them in four different groups (Harris et al. 2013) and by four different personalities in order to reach a higher type of granularity, based on their actual role within the company ("Archetype") and on personal features ("Personality"—according to the *Keirsey Temperament Sorter*). Identifying correctly the personality type of a data scientist is crucial to amplify his internal contribution and efficiency, as well as to maximize the resources employed to recruit him (Table 7.1).

In the table, a full disentanglement of data scientists' types is provided. The color roughly represents the partition between three main skills they possess—based on the survey run by Harris et al. (2013)—that are mathematics-statistic-modeling skills (blue), business ones (green), and coding abilities (red). Having this clear

Table 7.1 Data scientists' personality assessment and classification

Archetype/Personality	Artisan	Rational	Guardian	Idealist
Technical	*Gardener*: Data munging and coding	*Wrangler*: Algorithms implementation	*Architect*: System architecture and infrastructure	*Evangelist*: Enhancing technical community
Researcher	*Alchemist*: Experiments, exploration and ideas generation	*Groundbreaker*: Innovative methodologies and modeling	*Cruncher*: Analytical model optimization	*Champion*: Mentoring and teaching
Creative	*Trailblazer*: Spotting out new hidden connections	*Warlock*: Using new tools for new applications	*Catalyst*: Customer intelligence	*Visionary*: Information diffusion to public
Strategist	*Babelian*: Data Interpreter	*Fisherman*: Blue ocean strategy and monetization	*Mastermind*: Project management	*Advocate*: Promoting to management
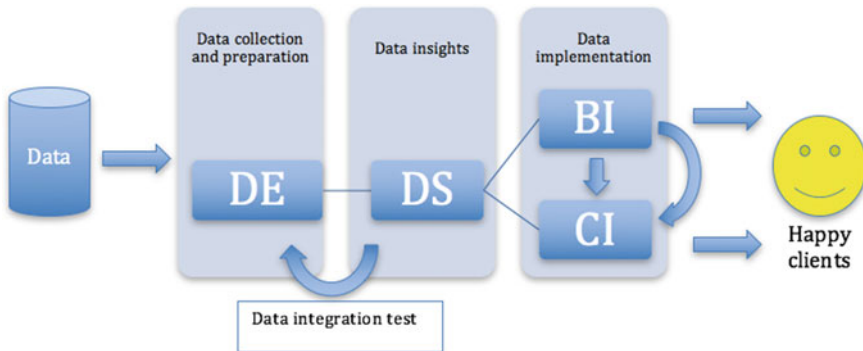
Fig. 7.2 Data science value chain

classification in mind may be argued to be a merely speculative and useless labeling exercise, but it is indeed extremely relevant because increases the data science team efficiency: identifying personal inclinations and aspirations would allocate the best people to the best job role, and common complaints and problems such the insufficient time for doing analysis, the poor data quality, and the excessive time spent in collecting and cleaning data (King and Magoulas 2015) would be eliminated—or better, they would be assigned to the right people. Furthermore, this framework would help identifying the minimum team structure to start with: on the main diagonal there are indeed the basic figures needed in order to properly establish a fully-functional data

science team. The *Gardener* (usually known also as data engineer) is in charge of maintaining the architecture and making the data available to the *Groundbreakers*, who are usually identified as proper data scientists, and that try to answer research questions and draw insights from data once they verified through tests that their models work. The insights are then passed to *Advocates* (business intelligence) and *Catalysts* (customer intelligence team), who respectively communicate that information to executives and use that to increase customers' satisfaction. The data process is illustrated in the following figure (Fig. 7.2).

Having these four different basic teams guarantees an efficient data-driven business and a sharp outcomes delivery, as soon as the communication across-teams and across-departments is well implemented. It is common practice to have short (five-minutes at most) stand-up internal meetings every morning to wrap up the daily objectives, works, and expected outcomes, as well as weekly meetings with other departments to align the work. Structuring the process as above proposed would finally increase the scalability of any data project.

References

Corea, F. (2016). Big data analytics: A management perspective. *Studies Series in Big Data, 21*. Heidelberg: Springer International Publishing.

Davenport, T. H., & Patil, D. J. (2012). Data scientist: The sexiest job of the 21st century. *Harvard Business Review, 90*(10), 70–76.

Harris, H. D., Murphy, S. P., & Vaisman, M. (2013). *Analyzing the analyzers*. O'Reilly Publishing.

Kandel, S., Paepcke, A., Hellerstein, J. M., & Heer, J. (2012). Enterprise Data analysis and visualization: An interview study. In *Proceedings of IEEE Visual Analytics Science & Technology (VAST)*.

King, J., & Magoulas, R. (2015). *2015 data science salary survey*. United States: O'Reilly Media Inc.

Liberatore, M., & Luo, W. (2012). ASP, the art and science of practice: A comparison of technical and soft skill requirements for analytics and or professionals. In *Interfaces 2013* (vol. 43, no. 2, pp. 194–197).

Chapter 8
AI and Speech Recognition

8.1 Conversation Interfaces

Conversational User Interfaces (CUI) are at the heart of the current wave of AI development. Although many applications and products out there are simply "*Mechanical Turks*"—which means machines that pretend to be automatized while a hidden person is actually doing all the work—there have been many interesting advancements in speech recognition from the symbolic or statistical learning approaches.

In particular, deep learning is drastically augmenting the abilities of the bots with respect to traditional NLP (i.e., bag-of-words clustering, TF-IDF, etc.) and is creating the concept of "*conversation-as-a-platform*", which is disrupting the apps market.

Our smartphone currently represents the most expensive area to be purchased per square centimeter (even more expensive than the square meters price of houses in Beverly Hills), and it is not hard to envision that having a bot as unique interfaces will make this area worth almost zero.

None of these would be possible though without heavily investing in speech recognition research. Deep Reinforcement Learning (DFL) has been the boss in town for the past few years and it has been fed by human feedbacks. However, I personally believe that soon we will move toward a B2B (bot-to-bot) training for a very simple reason: ***the reward structure***. Humans spend time training their bots if they are enough compensated for their effort.

This is not a new concept, and it is something Li Deng (Microsoft) and his group are really aware of. He actually provides a great threefold classification of AI bots:

– Bots that look around for information;
– Bots that look around for information to complete a specific task;
– Bots with social abilities and tasks (which he names *social bots* or *chatbots*).

For the first two, the reward structure is indeed pretty easy to be defined, while the third one is more complex, which makes it more difficult to be approached nowadays.

When this third class will be fully implemented, though, we would find ourselves living in a world where machines communicate among themselves and with humans in the same way. In this world, the ***bot-to-bot*** business model will be something ordinary and it is going to be populated by two types of bots: ***master bots*** and ***follower bots***.

I believe that research in speech recognition adds up, as well as the technology stacks in this specific space. This would result in some players creating "universal" bots (master bots) which everyone else will use as gateways for their (peripheral) interfaces and applications. The good thing of this centralized (and almost monopolistic) scenario is, however, that in spite of the two-levels degree of complexity, we won't have the black box issue affecting the deep learning movement today because bots (either master or follower) **will communicate between themselves in plain English rather than in any programming language**.

8.2   The Challenges Toward Master Bots

Traditionally, we can think of deep learning models for speech recognition as either ***retrieval-based*** models or ***generative-models***. The first class of models uses heuristics to draw answers from predefined responses given some inputs and context, while the latter generates new responses from scratch each time.

The state-of-art of speech recognition today has raised a lot since 2012, with deep-q networks (DQNs), deep belief networks (DBN), long short-term memory RNN, Gated Recurrent Unit (GRU), Sequence-to-sequence Learning (Sutskever et al. 2014), and Tensor Product Representations (for a great overview on speech recognition, look at Deng and Li 2013).

So, if DFL breakthroughs were able to improve our understanding of the *machine cognition*, what is preventing us from realizing the perfect social bots? Well, there are at least a couple of things I can think of.

First of all, **machine translation** is still in its infancy. Google has recently created a "Neural Machine Translation", a relevant leap ahead in the field, with the new version even enabling ***zero-short translation*** (in languages which they were not trained for).

Second, speech recognition is still mainly a supervised process. We might need to put further effort into Unsupervised Learning, and eventually even better integrate the symbolic and neural representations.

Furthermore, there are many nuances of human speech recognition which we are not able to fully embed into a machine yet. MetaMind is doing a great work in the space and it recently introduced **Joint Many-Tasks** (JMT) and the **Dynamic Coattention Network** (DCN), respectively an **end-to-end trainable model** which allows collaboration between different layers and a network that reads through

documents having *an internal representation of the documents conditioned on the question that it is trying to answer*.

Finally, the automatic speech recognition (ASR) engines created so far were either lacking *personality* or completely *missing the spatiotemporal context*. These are two essential aspects for a general CUI, and only a few works have been tried up to date (Yao et al. 2015; Li et al. 2016).

8.3   How Is the Market Distributed?

This was not originally intended to part of this chapter, but I found useful to go quickly through main players in the space in order to understand the importance of speech recognition in business contexts.

The history of bots goes back to Eliza (1966, the first bot ever), Parry (1968) to eventually ALICE and Clever in the nineties and Microsoft Xiaoice more recently, but it evolved a lot over the last 2–3 years.

I like to think about this market according to this two by two matrix. You can indeed classify bots as native or enablers, designed for either specific or generic applications. The edges of this classification are only roughed out and you might actually have companies operating at the intersection between two of these quadrants (Fig. 8.1):
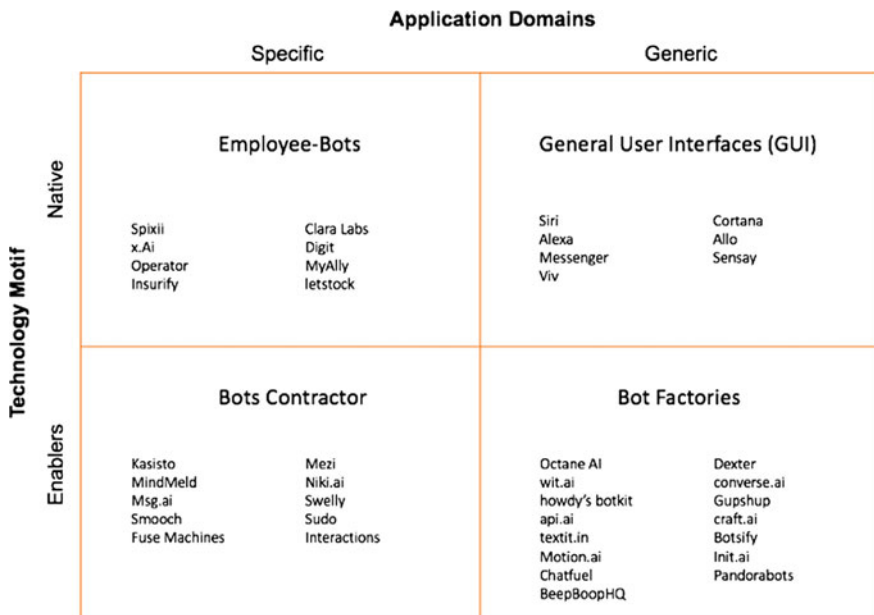
Fig. 8.1 Bots classification matrix

Following this classification, we can identify four different types of startups:

– **Employee-Bots**: these are bots that have been created within a specific industry or areas of application. They are stand-alone frameworks that do not necessitate extra training but are ready to plug and play;
– **General User Interfaces**: these are native applications that represent the purest aspiration to a general conversational interface;
– **Bots Contractors**: bots that are "hired" to complete specific purposes, but that were created as generalists. Usually cheaper and less specialized than their Employee brothers, live in a sort of symbiotic way with the parent application. It could be useful to think about this class as functional bots rather than industry experts (first class);
– **Bots Factories**: startups that facilitate the creation of your own bot.

A few (non-exhaustive) examples of companies operating in each group have been provided, but it is clear how this market is becoming crowded and really profitable.

8.4   Final Food for Thoughts

It is an exciting time to be working on deep learning for speech recognition. Not only the research community but the market as well are quickly recognizing the importance of the field as an essential step to the development of an AGI.

The current state of ASR and bots reflect very well the distinction between narrow AI and general intelligence, and I believe we should carefully manage the expectations of both investors and customers. I am also convinced is not a space in which everyone will have a slice of the pie and that a few players will eat most of the market, but it is so quick-moving that is really hard to make predictions on it.

References

Deng, L., & Li, X. (2013). Machine learning paradigms for speech recognition: An overview. In *IEEE Transaction on Audio, Speech, and Language Processing* 21(5).
Li, J., Galley, M., Brockett, C., Spithourakis, G., Gao, J., & Dolan, W. B. (2016). A persona-based neural conversation model. ACL (1).
Sutskever, I., Vinyals, O., & Le, Q. (2014). Sequence to sequence learning with neural networks. *NIPS, 2014,* 3104–3112.
Yao, K., Zweig, G., & Peng, B. (2015). Attention with intention for a neural network conversation model. CoRR abs/1510.08565.

Chapter 9
AI and Insurance

9.1 A Bit of Background

The insurance sector is one of the most old-fashioned and resistant-to-change space, and this is why AI will have a greater impact on that with respect to more receptive industries. The collection of data of new types (i.e., unstructured data such as reports, images, contracts, etc.) and the use of new algorithms are disrupting the sector in several ways.

Traditionally, an insurance company followed this type of process:

– Identifying pool of customers whom might be risk-assessed;
– Targeting those customers and assessing the risk for each class;
– Selling differently priced policies spreading the risks over the pool of customers;
– Try to retain those customers as long as possible offering lower price for longer contracts.

This is a really simplistic representation of the insurance business in the last fifty years, and I am aware that insurance experts might disagree with me in many different ways. There are a couple of further features to be pointed out: first of all, insurance has historically been *sold not bought*, which means that brokers and agents were essential to tracking new customers and to even retain old ones. In addition, it is an industry which is by definition rich of data because they collected anything they could, but is also one of the less advanced because either many of those data are unstructured or semi-structured, or the model used are quite old and simple.

Most of those data were easy to obtain because they were required to correctly price the coverage, while additional complimentary data were provided only by good customers who had incentives in providing as much data as possible to get a cheaper policy. Of course, this works the other way for bad customers, and this is a perspective on the phenomenon of "*adverse selection*" (i.e., bad customers are going to ask an insurance because they feel they will need it).

The adverse selection issue is though only one of the intrinsic challenges of the sector: **strong regulation**, high level of **fraud** attempts, and **complexity** are other features any incumbents should take care of. It is interesting to notice though that some of those are also specific barriers to entry for startups: they might attract indeed people who normally can get affordable insurance with a bigger competitor (adverse selection) and they usually have the capabilities for breaking down the risk complexity but not to support the funding need for risk coverages (so they need to work with incumbents rather than trying to replace them).

In spite of those problems, in the last decade, we noticed a new trend emerging. Insurances, in the effort of trying to reduce *moral hazard* problems, they started offering *premium discounts to their final customers* in order to get extra information. This occurred either through a questionnaire (asking **directly** the customer for further data in exchange for a lower price) or **indirectly** through devices (healthy devices, black boxes, etc.). The real issue though has been the engagement side of this proposal, because of the opposite nature of information, rewards, and human nature. The rewards offered were indeed either temporary or provided only once and people got lazy very quickly, while the information stream needed to be constant.

The following step has been the introduction of apps to let customers monitor by themselves their own data and behavior, sometimes even given away for free the device itself. Leaving the customer with full power on his data had though an inverse effect, because people did not have the motivation in tracking down their improvements, and they got upset at the same time because they felt they were not getting the most out of that opportunity.

Regardless of the specific innovative way in which insurers engaged customers, the process used in the insurance business did not change much in the past century. **Expert systems** and **knowledge engineering** dominated the sector setting the rules to be followed in internal workflows, but this is slowly changing with intelligent automation systems. We are actually migrating from rule-based decision systems to statistical learning and eventually machine learning.

9.2   So How Can AI Help the Insurance Industry?

AI is helping (or disrupting, depending on how you see the matter) the sector in different ways. First of all, it can help **increasing the customer engagement and retention** problem which has been just mentioned. The abundance of data can be used indeed to refine the customers' segmentation and provide personalized offers based on personal features. It also helps in **reducing the costs** through smart automatization or RPA (robotic process automation).

Second, AI is making people **more aware of the risks as well as habits**, and it is driving them toward better behaviors.

Furthermore, the better pricing and risk assessment that AI is introducing analyzing more granular data will make some people **uninsurable** (i.e., too risky to be

fairly priced and covered) as well as to turn back some previously uninsurable people into insurable customers again. The governments or central regulatory agencies should then start thinking about a "*pricing/risk threshold*" in which they intervene subsidizing the cost of relevant insurances (e.g., basic health coverage) in order to "*guarantee the uninsurables*".

Finally, it might be useful to think in terms of what an insurable risk is in order to see how AI can help with that.

According to Jin Park (Assistant Professor at IWU), an insurable risk is identifiable through the following five conditions:

– Large number of similar exposure units (mutuality);
– Accidental and unintentional loss (not predictable and independent from the insured customers);
– Determinable and measurable loss;
– Calculable chance of (not catastrophic/systemic) loss;
– Economically feasible premium.

AI is going to affect all those features: with a better and more detailed customer profiling, we won't need indeed to have such a large base of insured units. It will turn some frequent events into accidental (e.g., affecting drivers' behavior it will reduce the basic accidents into rare events) and it will improve our ability to forecast and compute both the probability and magnitude potential losses even in those cases too hard to be managed before. All the previous improvements will make many more premium under budgets, and therefore the conclusion is that AI will "*lower*" the threshold of what we consider nowadays an insurable risk, and it will make then more risks insurable.

9.3 Who Are the Sector Innovators?

There are plenty of startups out there working at the intersection of AI and insurance, and it essential to look at least at some of them to understand the future direction of the industry, as well as the kind of improvements AI is having in the insurtech space. An interesting thing to notice is that most of the innovation is happening in the UK rather than other countries, in all the segments proposed below.

Claim processing: Shift Technology skims the valid claims from the ones that deserve further validations; Tractable instead is trying to automatize experts task for insurances; ControlExpert has a specific focus on car claims; Cognotekt optimizes internal business processes, as well as Snapsheet does; Motionscloud offers instead mobile claim management solutions; and finally RightIndem aims to help insurers to deliver on-premise smoothing the claiming flow.

Virtual Agents and Chatbots: Spixii is an automated insurance agent who helps you buying any insurance coverage you might want; Cognicor is a virtual

assistant that offers customer care services; Conversica identifies which leads intend to purchase, while Your.MD is a personal health assistant that analyzes symptoms and produces pieces of advice. MedWhat instead uses EMR (medical records) to assist the patient as it was a virtual doctor, and Babylon gives medical advice taking care of tight budget constraints. Insurify is another personal insurance agent who works as a comparator for car insurances. What today is called simply chatbot is going to be renamed in a few years **robo-insurer**. There are already few examples of companies toward that goal: Risk Genius is indeed an intelligent comparator which identifies gaps in coverage for the customer and PolicyGenius looks for the best solution that fits customer's needs and characteristics, while Drive Spotter implements real-time video analytics to keep drivers safe(r). More generally, robo-insurers will be a quite wide class of agents who will end up providing different services, all of them with the final goal of helping the clients to undertake risk-mitigating actions and only cover the real (residual) risks.

Customers engagement: Oscar is probably the most successful insurtech company out there, with the final goal of making insurance simple and accessible to everyone through a great UX. Similar to Oscar is somehow Stride Health, while Brolly is a tool that helps customers in understanding their own needs and facilitates in one place all the insurance coverages in place, in a similar fashion to Knip. Adtelligence instead creates personalized offers and relevant products based on customer's characteristics. Captricity uses machine learning to convert hand-written files into structured data, and this can be used to better understand the final customer. Finally, ValChoice ranks the service of insurers to the benefit of the client.

Telematics: connected cars and telematics is a pretty big area itself, but it would be worthy to point out the work that Greenroad, Vnomics, and Telogis are doing in capturing driving behaviors and habits as well as computing fuel efficiency. Cambridge Mobile Telematics works similarly, although it uses smartphone data and mobile devices habits. Navdy is trying to revolutionizing the UI/UX within vehicles, displaying information in such a way that the driver does not get distracted. Lytx uses vision technology to provide real-time feedbacks to the driver.

Underwriting: AI can be (and actually is) used to spot out hidden correlations to granularly segment customers and risks in a more efficient way. Even though it might in theory possible to identify some algos that could perform better than others (see the work Wipro did for fraud detection), data always come first, at least for the next close future. Many companies operate in the space, as for instance Carpe Data that provides predictive algorithms and data products for property and casualty and life insurances through the analysis of public data (e.g., social media data). Atidot created a machine learning risk management platform, while Tyche uses unstructured data to optimize the underwriting and claims process. Big Cloud Analytics collects data from wearables and formulates health scores for a better risk assessment, while Cape Analytics uses computer vision techniques on geospatial data to improve the level of detail on current houses conditions. Dreamquark creates a more accurate representation of the medical datasets to be used for underwriting purposes by insurances, similarly to FitSense that offers also apps products. Melody

Health Insurance provides also low-cost insurances, while Uvamo uses AI to assess the risk of policy applications. A more accurate underwriting can even translate into covering events that are today quite risky (e.g., as MeteoProtect and Praedicat, and are doing for weather risk management).

Finally, on a side, it is worthy to point out to pure technological enablers as Instanda, which offers a management tool to the insurance providers to manage effectively and timely new products launched; Insly, a cloud-based platform for insurance brokers; and finally, SimpleInsurance is instead an e-commerce provider for product insurances.

P2P insurance: Lemonade, Friendsurance, and Guevara are peer-to-peer insurance startups focusing respectively on property and casualty insurance the first two, and car insurance the latter one.

Insurchain and Smart Contracts: these are companies in the insurance sector that are driven by *blockchain technology*. Elliptic offers real-time AML for bitcoin specifically, while Everledger is a permanent immutable ledger for diamond certification. Luther Systems is instead a stealth-mode company working on the standardization of smart contracts. Dynamis provides a P2P supplementary unemployment insurance product, while Saldo.mx provides micro-insurance policies on the blockchain. SafeShare covers multiple parties with insurance cover at short notice and for varying periods, and finally, Teambrella is another P2P insurance platform run on the blockchain.

Insurance on-demand: this class of startups put in customers' hand the entire insurance buying process. Trov is probably the best example of this new class of players and it allows to ensure things by simply taking a picture of them. Cuvva is quite similar but with a focus on car insurance, Sure and Airsurety on travel policies, and Back me up is another example of on-demand insurance. But this class does not include only the proper on-demand business model, but also insurance startups which provide products that vary by location, time, use, or customer. In other words, pay-per-mile business model (Metromile), micro-insurance policies (Neosurance), or eventually Insurance-as-a-service models (Digital Risks).

9.4   Concluding Thoughts

Yan identifies four elements which constitute the insurance profit structure: premium earned and the investment income from one hand, and underwriting cost and claim expenses from the other. AI is and will be able to improve the cost structure, increasing at the same time the competitiveness and enlarging the customer base accessible to insurers, while optimizing internal processes and enhancing the transparency and robustness of the compliance flow.

The greatest challenge I still see in insurance is the ***cultural mindset*** which might prevent insurance to adopt early AI solutions, although this won't probably have a long life given the incredible pressure to innovate the insurance providers are undergoing through.

Chapter 10
AI and Financial Services

10.1 Financial Innovation: Lots of Talk, Little Action?

The financial sector is historically one of the most resistant to change you might think of. It is then inevitable that big banks from one hand and startups from the other hand are creating a huge break in the financial industry and I believe this is happening not because of the use of a specific technology but rather because of their intrinsic cultural differences, diverse structural rigidity, and alternative cost-effective business models.

In other words, banks do not innovate either because they are too big to quickly adapt and follow external incentives or because they don't know how (and want to) truly change. This is not simply true in the industry but also in academia, where until the mid-nineties there were no relevant contributions to financial innovation at all (Frame and White 2002). In fact, in few survey articles (Cohen and Levin 1989; Cohen 1995) with more than 600 different articles and books quoted, **none of them was related to financial innovation subjects**.

Of course, things changed over the last five years, but my opinion is that was really out of necessity rather than a voluntary push-approach from the banking sector.

Financial innovation is, therefore, something which seems to be usually ***imported*** rather than ***internally generated***, and often more characterized by a ***product-innovation*** rather than a ***process one*** (although this might be a controversial opinion, I guess). Given the new technological paradigm (which is tightening the inner strong causal relationship between innovation and growth) it seems natural to wonder whether a better innovation model can be therefore imported by a different (and more successful) sector.

I found that there is a very specific and interesting case of a sector which had to '***innovate-to-survive***' rather than '***innovate-to-grow***': the biopharma industry (Baker 2003; Gans and Stern 2003; Fuchs and Krauss 2003; Lichtenthaler 2008).
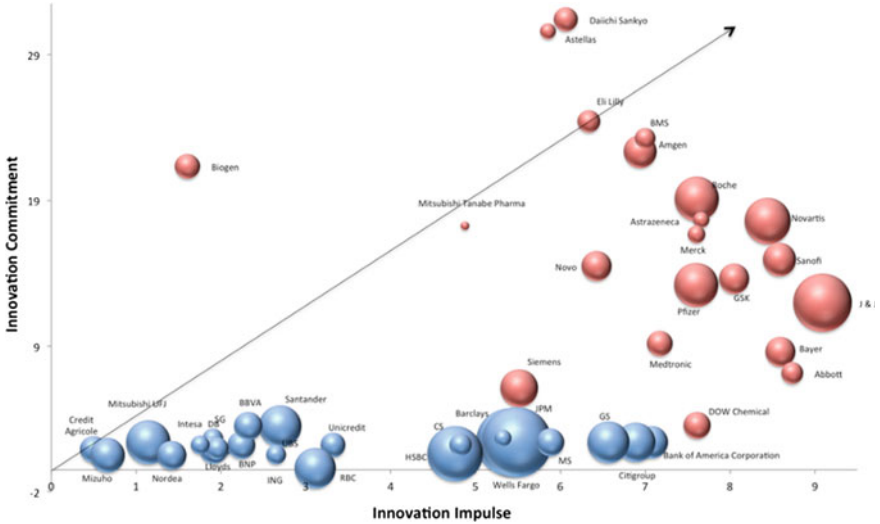
Fig. 10.1 Innovation transfer: the biopharma industry

Figure 10.1 **Innovation drivers map** (Corea 2015). **The biopharma companies feel more the urgency to innovate and are also more committed to that**. The graph I built plots 25 major banks (blue) and 25 major pharmaceutical (red) companies based on their Innovation Impulse and Commitment. The Impulse variable has been built using the number of patents a company filed (a proxy for the external pressure to innovate) and the number of recorded shareholders (a proxy for the internal pressure to innovate). The Commitment shows instead the R&D intensity (net sales) while the size of the bubbles the net income for each company. The data points were obtained by Medtrack, Osiris, and Zephyr in 2014.

10.2 Innovation Transfer: The Biopharma Industry

The biopharma industry is not a single sector but actually two different ones: the **biotech space**, populated by smaller companies that drive the research and exploration phase, and the **pharmaceutical companies**, big giants that through the last century became huge go-to-market and sales enterprises.

Hence, there is **pure (risky) innovation** from one hand and **pure commercialization skills** from the other…Is it something that we have already seen somewhere, didn't we? The biopharma industry and the financial sector suffer indeed from a strong polarized innovation.

What characterizes the industry is that the risky activity lies in the initial development process rather than in the market phase. The problem is not to match customer demand or find a market for your product, but it is actually developing the

molecule in the first place. The probability of success is extremely low and the timeline very long (10–15 years) and the 20-years patents give you only a temporary advantage. More importantly, it looks that only 3 out of 10 of the drugs produced are able to repay the development costs (Meyer 2002) and that most of the companies operate at loss while the top 3% companies alone generate almost 80% of the entire industry profits (Li and Halal 2002). A tough business, isn't it?

The biopharma industry is then no longer simply a ***human-intensive business*** but also a ***capital-demanding one***. Innovation is not ancillary but it is the quintessential driver to survive. And this is also why they had to identify a range of different methods to foster their growth-by-innovation: R&D, competitive collaboration schemes, venture funding, co-venture creation, built-to-buy deals, limited partnership agreements, etc.

It should be clear by now where I am heading to: **the financial industry doesn't strongly feel the need to innovate** as the biopharma sector and it is **not experimenting and pushing to create new models that might spread their innovation risk** and make it profitable.

10.3   Introducing AI, Your Personal Financial Disruptor

By now you might object "All good man, but financial services and biopharma are still so different, so why should I import innovation models from a sector which is completely different from mine?". Well, that's the catch: I don't think they are.

And the reason why they are becoming a lot more similar is precisely **Artificial Intelligence**.

AI is creating a strong pressure to innovate for the financial sector and has a **development cycle and characteristics** which are somehow similar to the biopharmaceutical one: it requires a long time to be created, implemented and correctly deployed (with respect to the financial industry standards, of course); it is highly technical and requires highly specialized talents; it is highly uncertain, because you need to experiment a lot before finding something that works; it is expensive, both in terms of time as well as monetary investments (talents, hardware, and data are really expensive); it is risky and the risk lies in the initial development phase, with a very high-payout but a high likelihood to fail as well.

But AI is also introducing a completely new speed and degree of trust in the financial industry, which lowers the tolerable mistakes at the same level of the biopharma sector. If your algorithms point out to the wrong product to sell or the wrong book to be recommended, it is not a big deal. If your system misinterprets some signals in the market or while developing a drug though, you end up **losing millions in seconds or even losing human lives**.

It is then not only stretching out issues that intrinsically belong to the financial sector such as **regulation** or **accountability**, but it is also bringing new problems such as **biased data or the lack of transparency** to the picture (specifically in consumer applications).

And last but not least, AI is making the question mark on the "***build vs buy***" matter bigger than even in FS, the same as it was in the biopharma industry back in the nineties and that culminated in the current biotech-pharmaceutical dichotomy (if you are wondering anyway, this choice is all focused around on your ***data capacity, team and project scalability***, and ***uniqueness of the project*** with respect to your competitors—do you have enough data to train an ANI? Can your team/project scale? Is the ANI unique or something your peers are doing or need to do as well?).

This is why I believe AI in financial services to be extremely important—not much for the specific innovation or product it is introducing but rather because ***it is revolutionizing a centuries-old industry innovation flow from the ground***.

10.4   Segmentation of AI in Fintech

Artificial Intelligence is using structured and unstructured data in financial services to improve the customer experience and engagement, to detect outliers and anomalies, to increase revenues, reduce costs, find predictability in patterns and increase forecasts reliability…but it is not so in any other industry? We all know this story, right? So what is really peculiar about AI in financial services?

First of all, financial services is an industry full of data. You might expect this data to be concentrated in big financial institutions' hands, but most of them are actually public and thanks to the new **EU payment directive (PSD2) larger datasets are available** to smaller players as well. AI can then be easily developed and applied because the barriers to entry are lower with respect to other sectors.

Second, many of the underlying processes can be relatively **easier to be automatized** while many others can be improved by either brute force computation or speed. And historically is one of the sectors that needed this type of innovation the most, is incredibly competitive and is always looking for some new source of ROI. **Bottom line: the marginal impact of AI is greater than in other sectors**.

Third, the **transfer of wealth across different generations** makes the field really fertile for AI. AI needs (a lot of) innovative data and above all feedback to improve, and millennials are not only happy to use AI as well as providing feedback, but apparently even less concerned about privacy and giving away their data.

There are also, of course, a series of **specific challenges for AI** in financial sector that limit a smooth and rapid implementation: legacy systems that do not talk to each other; data silos; poor data quality control; lack of expertise; lack of management vision; lack of cultural mindset to adopt this technology.

So what is missing now is only having an overview of the AI fintech landscape. There are also plenty of maps and classification of AI fintech startups out there, so I am not introducing anything new here but rather simply giving you my personal framework:

– **Financial Wellness**: this category is about making the end-client life better and easier and it includes ***personalized financial services***; ***credit scoring***; automated financial advisors and planners that assist the users in making financial decisions (***robo-advisor, virtual assistants, and chatbots***); smart wallets that coach users differently based on their habits and needs. *Examples include [robo-advisors and conversational interfaces] Kasisto; Trim; Penny; Cleo; Acorns; Finge nius; Wealthfront; SigFig; Betterment; LearnVest; Jemstep; [credit scoring] Aire; TypeScore; CreditVidya; ZestFinance; Applied Data Finance; Wecash;*

– **Blockchain**: I think that, given the importance of this instrument, it deserves a separate category regardless of the specific application is being used for (which may be payments, compliance, trading, etc.). *Examples include: Euklid; Paxos; Ripple; Digital Asset;*

– **Financial Security**: this can be divided into **identification** (payment security and physical identification—**biometrics and KYC**) and **detection** (looking for fraudulent and abnormal financial behaviour—**AML and fraud detection**). *Examples include, respectively: EyeVerify; Bionym; FaceFirst; Onfido; and Feedzai; Kount, APEX Analytics;*

– **Money Transfer**: this category includes payments, peer-to-peer lending, and debt collection. *Examples include: TrueAccord; LendUp; Kabbage; LendingClub;*

– **Capital Markets**: this is a big section, and I tend to divide it into five main subsections:

  • **Trading** (either algo trading or trading/exchange platforms). *Examples include: Euclidean; Quantestein; Renaissance Technologies, Walnut Algori thms; EmmaAI; Aidyia; Binatix; Kimerick Technologies;Pit.ai;Sentient Techn ologies; Tickermachine; Walnut Algorithm; Clone Algo; Algoriz; Alpaca; Portfolio123; Sigopt;*

  • **Do-It-Yourself Funds** (either crowdsource funds or home-trading). *Examples include: Sentifi; Numerai; Quantopian; Quan tiacs; QuantConnect; Inovance;*

  • **Markets Intelligence** (information extraction or insights generation). *Examples include: Indico Data Solutions; Acuity Trading; Lucena Research; Dataminr; Alphasense; Kensho Technologies; Aylien; I Know First; Alpha Modus; ArtQuant;*

  • **Alternative Data** (most of the alternative data applications are in capital markets rather than broader financial sector so it makes sense to put it here). *Examples include: Cape Analytics; Metabiota; Eagle Alpha;*

  • **Risk Management** (this section is more a residual subcategory because most of the time startups in this group fall within other groups as well). *Examples include: Ablemarkets; Financial Network Analysis.*

10.5   Conclusions

I am arguing since the beginning of the article that AI is making financial services and biopharma much more alike, and that the FS industry might learn something from how the other industry innovates.

The reality is that the financial industry has also very specific traits and challenges it needs to overcome.

The biggest difference I currently see in that is the effect AI is having on the physical products market: while in almost any sector AI is used with the final goal of creating or improving new products (and this is true also for drug development, for example) in the financial ecosystem is having exactly the opposite effect. **AI is making the industry more digitalized than ever before**. Its final goal will be to create the (frictionless) bank of the future: no branches, no credit cards, no frauds, no menial reporting activities. A bank-as-a-platform with modular components that increases our financial literacy and has no physical products or spaces.

References

Baker, A. (2003). Biotechnology's growth-innovation paradox and the new model for success. *Journal of Commercial Biotechnology, 9*(4), 286–288.

Cohen, W. (1995). Empirical studies of innovative activity (Ch. 6). In P. Stoneman (Ed.), *Handbook of the economics of innovation and technological change* (pp. 182–264). Cambridge, Massachusetts: Blackwell.

Cohen, W., & Levin, R. (1989). Empirical studies of innovation and market structure (Ch. 18). In R. Schmalensee & R. Willig (Eds.), *Handbook of Industrial Organization* (Vol. 2, pp. 1059–1107). Amsterdam, North-Holland.

Corea, F. (2015). What finance can learn from biopharma industry: A transfer of innovation models. *Expert Journal of Finance, 3,* 45–53.

Frame, W. S., White, L. J. (2002). Empirical studies of financial innovation: Lots of talk, little action?. *Working paper. Federal Reserve Bank of Atlanta*, N. 2002–12.

Fuchs, G., & Krauss, G. (2003). Biotechnology in comparative perspective. In G. Fuchs (Ed.), *Biotechnology in Comparative Perspective* (pp. 1–13). New York: Routledge.

Gans, J., Stern, S. (2003). *Managing ideas: Commercialization strategies for biotec hnology. Intellectual Property Research Institute of Australia Working Paper 01/03,* 1–24.

Li, J., Halal, W. E. (2002). Reinventing the biotech manager. *Nature Biotechnology*, *20*(6s), 61–3.

Lichtenthaler, U. (2008). Open innovation in practice: An analysis of strategic approaches to technology transactions. *IEEE Transactions on Engineering Management, 55*(1), 148–157.

Meyer, F. J. (2002, November 25). Business models that biotech companies employ. *Enterprise Development KFBS Biotech Speakers Series*.

Chapter 11
AI and Blockchain

11.1 Non-technical Introduction to Blockchain

A blockchain is a **secure distributed immutable database shared by all parties in a distributed network** where transaction data can be recorded (either *on-chain* for basic information or *off-chain* in case of extra attachments) and easily audited.

Put simply (with Bank of England's words), the blockchain is "a *technology that allows people who don't know each other to trust a shared record of events*".

The data are stored in rigid structures called **blocks**, which are connected to each other in a **chain** through a *hash* (each block also includes a *timestamp* and a link to the previous block via its *hash*). The blocks have a header, which includes metadata, and a content, which includes the real transaction data. Since every block is connected to the previous one, as the number of participants and blocks grow, it is extremely hard to modify any information without having the network consensus.

The network can validate the transaction through different mechanisms, but mainly through either a "*proof-of-work*" or a "*proof-of-stake*". A **proof-of-work** (Nakamoto 2008) asks the participants (called "*miners*") to solve complex mathematical problems in order to add a block, which in turn require a ton of energy and hardware capacity to be decoded. A **proof-of-stake** (Vasin 2014) instead tries to solve this energy efficiency issue attributing (roughly) more mining power to participants who own more coins (there are many variations of it and some skepticism around its famous "*nothing at stake*" problem.

Additional mechanisms are the Byzantine-fault-tolerant algorithm (Castro and Liskov 2002), the Quorum slicing (Mazieres 2016), as well as variations of the Proof-of-stake (Mingxiao et al. 2017), but we will not get into those now.

The final characteristic that needs to be explained is the category of blockchain based on the different network access permission, i.e., whether it is free for anyone to view it (**permissionless vs. permissioned**) or to participate in the consensus formation (**public vs. private**). In the former case, anyone can access and read

or write data from the ledger, while in the latter one predetermined participants have the power to join the network (and of course only in the public permissionless case a reward structure for miners has been designed).

It should be clear by now the intrinsic power of this technology, which is not simply a disruptive innovation but rather a *foundational technology* that aims to "*change the scope of intermediation*" (Catalini and Gans 2017). Distributed ledger technologies will indeed reduce both the costs of verification and networking, influencing then the market structure and eventually allowing the creation of new marketplaces. Iansiti and Lakhani (2017) also drew abrilliant parallel between blockchain and TCP/IP in a recent work (which I highly recommend), showing how blockchain is slowly going through the four phases that identify previous foundational technologies such as the TCP/IP, i.e., single-use, localized use, substitution, and transformation. As they explained, the "*novelty*" of such a technology makes it harder for people to understand the solution domain, while its "*complexity*" requires a larger institutional change to foster an easy adoption.

However, it is also true that the blockchain is shifting the traditional business models distributing value in an opposite way with respect to previous stacks: if it made more sense to invest in applications rather than protocol technologies fifteen years ago, in a blockchain world the value is concentrated in the shared protocol layer and only marginally at the application level (see the "***Fat Protocol***" theory by Joel Monegro).

To conclude this introductory section, I will just mention on the fly the possibility for the blockchain to not simply allow for transactions but also the possibility to create **(smart) contracts** that are triggered by specific events and threshold and that are traceable and auditable without effort.

11.2   A Digression on Initial Coin Offerings (ICOs)

A big hype is nowadays surrounding this new phenomenon of the *Initial Coin Offerings* (*ICOs*). Even if many people are pouring money into that because of its resemblance to the most common (and valuable) Initial Public Offerings (IPOs), an ICO is nothing more than a **token sale**, where a token is the *smallest functional unit of a specific network* (or application).

ICOs experts (if any) will forgive my approximate definition, but an ICO is a hybrid concept that has elements of a ***shares allocation***, a ***pre-sales/ crowdfunding campaign***, and a ***currency*** with a limited power and application's domain.

It is definitely an interesting innovation that introduces new unregulated ways to raise capitals, but it also poses several issues to an unprepared community. I am happy to receive feedback on this, but I would distil the key points of an ICO evaluation in what follows:

– A token has an additional utility with respect to the exchange of value and companies selling token with the **only goal of raising capital are sending a bad signal** to the market. Tokens are needed to create a users' base and to incentivize stakeholders to participate in the ecosystem at the earliest stage. **A good white paper is not enough;**
– Be wary of token sales that are **uncapped**;
– Be wary of token sales that have **no time limit**;
– Be wary of token sales that do not clearly state the (present and future) **number** as well as the **value of the token** (it could sound absurd, but you may be surprised of how non-transparent an ICO can look like).

11.3 How AI Can Change Blockchain

Although extremely powerful, a blockchain has its own limitations as well. Some of them are technology-related while others come from the old-minded culture inherited from the financial services sector, but all of them can be affected by AI in a way or another:

– **Energy consumption**: *mining* is an incredibly hard task that requires a ton of energy (and then money) to be completed (O'Dwyer and Malone 2014). AI has already proven to be very efficient in optimizing energy consumption, so I believe similar results can be achieved for the blockchain as well. This would probably also result in lower investments in mining hardware;
– **Scalability**: the blockchain is growing at a steady pace of 1 MB every 10 min and it already adds up to 85 GB. Nakamoto (2008) first mentioned "*blockchain pruning*" (i.e., deleting unnecessary data about fully spent transactions in order to not hold the entire blockchain on a single laptop) as a possible solution but AI can introduce new decentralized learning systems such as **federated learning**, for example, or new data sharding techniques to make the system more efficient;
– **Security**: even if the blockchain is almost impossible to hack, its further layers and applications are not so secure (e.g., the DAO, Mt Gox, Bitfinex, etc.). The incredible progress made by machine learning in the last two years makes AI a fantastic ally for the blockchain to guarantee a secure applications deployment, especially given the fixed structure of the system;
– **Privacy**: the privacy issue of owning personal data raises regulatory and strategic concerns for competitive advantages (Unicredit 2016). *Homomorphic encryption* (performing operations directly on encrypted data), the **Enigma project** (Zyskind et al. 2015) or the **Zerocash project** (Sasson et al. 2014), are definitely potential solutions, but I see this problem as closely connected to the previous two, i.e., scalability and security, and I think they will go *pari passu*;
– **Efficiency**: Deloitte (2016) estimated the total running costs associated with validating and sharing transactions on the blockchain to be as much as

$600 million a year. An intelligent system might be eventually able to compute on the fly the likelihood for specific nodes to be the first performing a certain task, giving the possibility to other miners to shut down their efforts for that specific transaction and cut down the total costs. Furthermore, even if some structural constraints are present, a better efficiency and a lower energy consumption may reduce the ***network latency*** allowing then faster transactions;

– **Hardware**: miners (and not necessarily companies but also individuals) poured an incredible amount of money into specialized hardware components. Since energy consumption has always been a key issue, many solutions have been proposed and much more will be introduced in the future. As soon as the system becomes more efficient, some piece of hardware might be converted (sometimes partially) for neural nets use (the mining colossus Bitmain is doing exactly this);

– **Lack of talent**: this is leap of faith, but in the same way we are trying to automate data science itself (unsuccessfully, to my current knowledge), I don't see why we couldn't create virtual agents that can create new ledgers themselves (and even interact on it and maintain it);

– **Data gates**: in a future where all our data will be available on a blockchain and companies will be able to directly buy them from us, we will need help to grant access, track data usage, and generally make sense of what happens to our personal information at a computer speed. This is a job for (intelligent) machines.

11.4 How Blockchain Can Change AI

In the previous section, we quickly touched upon the effects that AI might eventually have on the blockchain. Now instead, we will make the opposite exercise understanding what impact can the blockchain have on the development of machine learning systems. More in details, blockchain could:

– **Help AI explaining itself (and making us believe it)**: the AI black-box suffers from an explainability problem. Having a clear audit trail can not only improve the **trustworthiness** of the data as well as of the models but also provide a clear route to trace back the machine decision process;

– **Increase AI effectiveness**: a secure data sharing means more data (and more training data), and then better models, better actions, better results…and better new data. Network effect is all that matter at the end of the day;

– **Lower the market barriers to entry**: let's go step by step. Blockchain technologies can secure your data. So why won't you store all your data privately and maybe sell it? Well, you probably will. So, first of all, blockchain will foster the creation of **cleaner and more organized personal data**. Second, it will allow the emergence of **new marketplaces**: a ***data marketplace*** (low-hanging fruit); a ***models marketplace*** (much more interesting); and finally even an ***AI marketplace*** (see what Ben Goertzel is trying to do

with SingularityNET). Hence, easy data-sharing and new marketplaces, jointly with blockchain data verification, will provide a more fluid integration that lowers the barrier to entry for smaller players and shrinks the competitive advantage of tech giants. In the effort of lowering the barriers to entry, we are then actually solving two problems, i.e., providing a *wider data access* and a more efficient *data monetization mechanism*;

– **Increase artificial trust**: as soon as part of our tasks will be managed by autonomous virtual agents, having a clear audit trail will help **bots to trust each other** (and us to trust them). It will also eventually increase every **machine-to-machine interaction** (Outlier Ventures 2017) and transaction providing a secure way to share data and coordinate decisions, as well as a robust mechanism to reach a quorum (extremely relevant for swarm robotics and multiple agents scenarios). Rob May expressed a similar concept in one of his last newsletters (that I highly recommend—you should definitely subscribe);

– **Reduce catastrophic risks scenario**: an AI coded in a DAO with specific smart contracts will be able to only perform those actions, and nothing more (it will have a limited action space then).

In spite of all the benefits that AI will receive from an interaction with blockchain technologies, I do have one big question with no answer whatsoever.

AI was born as in an open-source environment where data was the real moat. With this data democratization (and open-source software) how can we be sure that AI will prosper and will keep being developed? What would be the new moat? My only guess at the moment? Talent.

11.5 Decentralized Intelligent Companies

There are plenty of landscapes of blockchain and cryptocurrencies startups out there. I am anyway only interested in those companies working at the intersection (or the *convergence*, as someone calls it) of AI and blockchain, which apparently are not that many. They are mainly concentrated in San Francisco area and London, but there are examples in New York, Australia, China, as well as some European countries.

They are indeed so few of them that is quite hard to classify them into clusters. I usually like to try to understand the underlying patterns and the type of impact/application certain groups of companies are having in the industry, but in this case is extremely difficult given the low number of data points so I will simply categorize them as follows:

– **Decentralized Intelligence**: TraneAI (training AI in a decentralized way); Neureal (peer-to-peer AI supercomputing); SingularityNET (AI marketplace); Neuromation (synthetic datasets generation and algorithm training platform); AI Blockchain (multi-application intelligence); BurstIQ

(healthcare data marketplace); AtMatrix (decentralized bots); Open Minedproject (data marketplace to train machine learning locally); Synapse.ai (data and AI marketplace); Dopamine.ai (B2B AI monetization platform); Effect.ai (decentralized AI workforce and services marketplace);

– **Conversational Platform**: Green Running (home energy virtual assistant); Talla (chatbot); doc.ai (quantified biology and healthcare insights);
– **Prediction Platform**: Augur (collective intelligence); Sharpe Capital (crowd-source sentiment predictions);
– **Intellectual Property**: Loci.io (IP discovery and mining);
– **Data provenance**: KapeIQ (fraud detection on healthcare entities); Data Quarka (facts checking); Priops (data compliance); Signzy (KYC);
– **Trading**: Euklid (bitcoin investments); EthVentures (investments on digital tokens). For other (theoretical) applications in finance, see Lipton (2017);
– **Insurance**: Mutual.life (P2P insurance), Inari (general);
– **Miscellaneous**: Social Coin (citizens' reward systems); HealthyTail (pet analytics); Crowdz (e-commerce); DeepSee (media platform); ChainMind (cybersecurity).

A few general comments:

(i) It looks interesting that many AI-blockchain companies have **larger advisory board than teams**. It might be an early sign that the convergence is not fully realized yet and there are more things we don't understand that those ones we know;

(ii) I am personally very excited to see the development of the first category (**decentralized intelligence**) but I also see a huge development in **conversational** and **prediction platforms** as well as **intellectual property**. I grouped other examples under "miscellaneous" because I don't think at this stage they represent a specific category but rather only single attempt to match AI and blockchain;

(iii) **Those companies are incredibly hard to evaluate**. The websites are often cryptic enough to not really understand what they do and how (a bit paradoxical if you want to buy the blockchain transparency paradigm) and technology requires a high tech-education to be fully assessed. Cutting through the hype is a demanding task and this makes it very easy to be fooled. But let me give you a concrete example: **ever heard of Magos AI**? In the effort of researching companies for this post, I found myself reading several articles on this forecasting blockchain AI-driven platform company (Wired, Prnewswire, etc.), which just did an ICO for over half a million dollars and that made great promises on its deliverables. The website didn't work—weird, if you consider that they need to share material/information on the ICOs. But you know, it might happen. I made then an extra effort because I read it on Wired and I was curious to know more about it. I was able to find its co-founders, which I couldn't find eventually on Linkedin. Weird again. Well, there are people who do not like socials, fair enough, especially

if you consider that until three months ago there was no proof of the company existence whatsoever. Let me look into the rest of the team. Nothing even there, and no traceable indications of their previous experiences (except for the CTO master in analytics, that I found no proof of). I tried to then dig into the technology: white papers, blue papers, yellow papers, you name it. I only found reviews of them, no original copies. Final two steps: I don't consider myself an expert in blockchain at all, but I read, a lot. And I also believe I am fairly knowledgeable when it comes to AI and what is happening in the industry. These guys claimed they created 5 different neural nets that could achieve the same accuracy in complex different domains than Libratus (or DeepStack) reached in Poker, but I never heard of them—very weird. Well, you know what? Maybe I could write them and meet them to understand. Their address points to the AXA office in Zurich. Ah. After 5 min of research, I finally Google the two key words: "Magos scam". It seems these guys took the money and disappeared. They are surely building the 6 neural net somewhere, so stay tuned.

My point here is that exponential technologies are fantastic and can advance mankind, but as much as the benefits increase also the potential "*negative convergence*" increases exponentially. Stay alert.

11.6 Conclusion

Blockchain and AI are the two extreme sides of the technology spectrum: one fostering centralized intelligence on close data platforms, the other promoting decentralized applications in an open-data environment. However, if we find an intelligent way to make them working together, the total positive externalities could be amplified in a blink.

There are of course technical and ethical implications arising from the interaction between these two powerful technologies, as for example how do we edit (or even forget) data on a blockchain? Is an ***editable blockchain*** the solution? And is not an AI-blockchain pushing us to become data hoarder?

Honestly, I think the only thing we can do is keep experimenting.

References

Castro, M., & Liskov, B. (2002). Practical byzantine fault tolerance and proactive recovery. *ACM Transactions on Computer Systems, 20*(4), 398–461.

Catalini, C., & Gans, J. S. (2017). *Some simple economics of the blockchain*. MIT Sloan School Working Paper: 5191–16.

Deloitte. (2016). *Blockchain enigma. Paradox. Opportunity*. White Paper.

Iansiti, M., & Lakhani, K. R. (2017). *The truth about blockchain* (pp. 118–127). Harvard Business Review, January–February 2017.

Lipton, A. (2017). *Blockchains and distributed ledgers in retrospective and perspective*. arXiv:1703.01505.

Mazieres, D. (2016). *The stellar consensus protocol: A federated model for internet-level consensus*. White Paper.

Mingxiao, D., Xiaofeng, M., Zhe, Z., Xiangwei, W., & Qijun, C. (2017). A review on consensus algorithm of blockchain. In *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC) Banff Center*, Banff, Canada, October 5–8, 2017.

Nakamoto, S. (2008). *Bitcoin: A peer-to-peer electronic cash system*. White Paper.

O'Dwyer, K. J., & Malone, D. (2014). Bitcoin mining and its energy footprint. In *25th IET Irish Signals and Systems Conference 2014 and 2014 China-Ireland International Conference on Information and Communications Technologies (ISSC 2014/CIICT 2014)*, Limerick (pp. 280–285).

Outlier Ventures. (2017). *Blockchain-enabled convergence*. White Paper.

Sasson, E. B., Chiesa, A., Garman, C., Green, M., Miers, I., Tromer, E., & Virza, M. (2014). Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE symposium on security and privacy* (SP) (pp. 459–474).

Unicredit. (2016). *Blockchain technology and applications from a financial perspective*. Technical Report.

Vasin, P. (2014). *BlackCoin's proof-of-stake protocol v2*. White Paper.

Zyskind, G., Nathan, O., & Pentland, A. (2015). *Enigma: Decentralized computation platform with guaranteed privacy*. arXiv:1506.03471.

Chapter 12
New Roles in AI

12.1 Hiring New Figures to Lead the Data Revolution

It has been said that this new wave of exponential technologies will threaten a lot of jobs, both blue and white-collar ones. But if from one hand many roles will disappear, from the other hand in the very short-term we are observing new people coming out from the crowd to lead this revolution and set the pace.

These are the people who really understand both the technicalities of the problems as well as have a clear view of the business implications of the new technologies and can easily plan how to embed those new capabilities in enterprise contexts.

Hence, I am going to briefly present three of them, i.e., the **Chief Data Officer** (CDO), the **Chief Artificial Intelligence Officer** (CAIO) and the **Chief Robotics Officer** (CRO). Sad to be said, I never heard about a '*Chief of Data Science*', but for some strange reasons, the role is usually called either '*Head of Data Science*' or '*Chief Analytics Officer*' (as if data scientist won't deserve someone at C-level to lead their efforts).

Let's see then who they are and what they would be useful for.

12.2 The Chief Data Officer (CDO)

Apparently, it is a new role born in a lighter form straight after the financial crisis springing from the need to have a central figure to deal with technology, regulation and reporting.

Therefore, the CDO is basically the guy who acts as a **liaison between the CTO (tech guy) and the CAO/Head of Data Science (data guy)** and takes care of data quality and data management.

Actually, her final goal is **to guarantee that everyone can get access to the right data in virtually no time**.

In that sense, a CDO is the guy in charge of '**democratizing data'** within the company.

It is not a static role, and it evolved from simply being a *facilitator* to being a *data governor*, with the tasks of defining data management policies and business priorities, shaping not only the data strategy, but also the frameworks, procedures, and tools. In other words, he is a kind of 'Chief of Data Engineers' (if we agree on the distinctions between data scientists, who actually deal with modeling, and data engineers, who deal with data preparation and data flow).

> The difference between a CIO and CDO (apart from the words data and information…) is best described using the bucket and water analogy. The CIO is responsible for the bucket, ensuring that it is complete without any holes in it, the bucket is the right size with just a little bit of spare room but not too much and its all in a safe place. The CDO is responsible for the liquid you put in the bucket, ensuring that it is the right liquid, the right amount and that's not contaminated. The CDO is also responsible for what happens to the liquid, and making the clean vital liquid is available for the business to slake its thirst. (Caroline Carruthers, Chief Data Officer Network Rail, and Peter Jackson, Head of Data Southern Water)

Interestingly enough, the role of the CDO as we described it is both *vertical* and *horizontal*. It spans indeed across the entire organization even though the CDO still needs to report to someone else in the organizational chart. Who the CDO reports to will be largely determined by the organization he is operating in. Furthermore, it is also relevant to highlight that a CDO can be found more likely in larger organizations rather than small startups. The latter type is indeed usually set up to be data-driven (with a *forward-looking approach*) and therefore the CDO function is already embedded in the role who designs the technological infrastructure/data pipeline.

It is also true that not every company has a CDO, so how do you decide to eventually get one? Well, simply out of internal necessity, strict incoming regulation, and because all your business intelligence projects are failing because of data issues. If you have any of these problems, you might need someone who pushes the "fail-fast" principle as the data approach to be adopted throughout the entire organization, who considers data as a company asset and wants to set the fundamentals to allow fast trial and error experimentations. And above all, someone **who is centrally liable and accountable for anything about data**.

A CDO is then the **end-to-end data workflow responsible and it oversees the entire data value chain.**

Finally, if the CDO will do his job in a proper way, you'll be able to see two different outcomes: first of all, the board will stop asking for quality data and will have clear in mind what every team is doing. Second, and most important, a **good CDO aims to create an organization where a CDO has no reasons to exist**. It is counterintuitive, but basically, a CDO will do a great job when the company won't need a CDO anymore because **every line of business will be responsible and liable for their own data**.

In order to reach his final goal, he needs to prove from the beginning that not investing in higher data quality and frictionless data transfer might be a source of inefficiency in business operations, resulting in non-optimized IT operations and making compliance as well as analytics much less effective.

12.3 The Chief Artificial Intelligence Officer (CAIO)

If the CDO is somehow an already consolidated role, the CAIO is nothing more than a mere industry hypothesis (not sure I have seen one yet, although the strong ongoing discussions between AI experts and sector players—see here and here for two opposite views on the topic). Moreover, the creation of this new role highlights the emergence of two different schools of thought of enterprise AI, i.e., **centralized versus decentralized AI implementation**, and a clear cost-benefit analysis to understand which approach will work better is still missing.

My two cents are that elevating AI to be represented at the board level means to really become an AI-driven company and embed AI into every product and process within your organization—and I bet not everyone is ready for that.

So, let's try to sketch at a glance the most common themes to consider when talking about a CAIO:

– **Responsibilities** (*what he does*): a CAIO is someone who should be able to connect the dots and **apply AI across data and functional silos** (this is Andrew Ng's view, by the way). If you also want to have a deeper look at what a CAIO job description would look like, check out here the article by Tarun Gangwani;
– **Relevance** (*should you hire a CAIO?*): you only need to do it if you understand that **AI is no longer a competitive advantage to your business but rather a part of your core product** and business processes;
– **Skills** (*how do you pick the right guy?*): first and more important, a CAIO has to be a 'guiding light' within the AI community because he will be one of your decisive assets to **win the AI talent war**. This means that he needs to be highly respected and trusted, which is something that comes only with a strong understanding of **foundational technologies and data infrastructure**. Finally, being a cross-function activity, he needs to have the right balance between **willingness to risk and experiment to foster innovation** and **attention to product and company needs** (he needs to support different lines of business);
– **Risks** (*is a smart move hiring a CAIO?*): there are two main risks, which are (i) the misalignment between technology and business focus (you tend to put more attention on technology rather than business needs), and (ii) every problem will be tackled with AI tools, which might not be that efficient (this type of guys are super trained and will be highly paid, so it is natural they will try to apply AI to everything).

Where do I stand on that? Well, my view is that a CAIO is something which makes sense, **even though only temporarily**. It is an essential position to allow a smooth transition for companies who strive for becoming AI-driven firms, but I don't see the role to be any different from what a smart tech CEO of the future should do (of course, supported by the right lower management team). However, for the next decade having a centralized function with the task of **using AI to support the business lines** (50% of the time) and **foster innovation internally** (50% of the time) it sounds extremely appealing to me.

In spite of all the predictions I can make, the reality is that the relevance of a CAIO will be determined by how we will end up approaching AI, i.e., whether it will be eventually considered a mere instrument **(AI-as-a-tool)** or rather a proper business unit **(AI-as-a-function).**

12.4   The Chief Robotics Officer (CRO)

We moved from the CDO role, which has been around for a few years now, to the CAIO one, which is close to being embedded in organizational charts. But the Chief Robotics Officer is a completely different story.

Even if someone is speaking about the importance of it, it is really not clear what his tasks would be and what kind of benefits would bring to a company, and envisaging this role requires a huge leap of imagination and optimism about the future of work (and business).

In few words, what a CRO will be supposed to take care of is **managing the automated workforce of the company**. To use Gartner's words, '*he will oversee the blending of human and robotic workers'*. He will be responsible of the overall automatization of workflows and to integrate them smoothly into the normal design process and daily activities.

I am not sure I get the importance of this holistic approach to enterprise automation, although I recognize the relevance of having a central figure who will actively keep track and communicate to employees all the changes made in transforming a manual activity/process into an automated one.

Another interesting point is who the CRO will report to, which is of course shaped by his real functions and goals. If robotics is deeply routed into the company and allows to create or access new markets, **a CRO might directly report to the CEO**. If his goal is instead to automatize internal processes to achieve a higher efficiency, **he will likely report to the COO** or to a strategic CxO (varying on industry and vertical).

My hypothesis is that this is going to be a **strategic role** (and not a technical one, as you might infer from the name) which, as the CAIO, might have a positive impact in the short term (**especially in managing the costs of adopting early robotics technologies**) but no reason to exist in the longer term. It is easier to think

about it in physical product industries rather than digital products or services companies, but automation will likely happen in a faster way in the latter, so we will end up having a *Chief of Physical Robotics Officer* (to manage the supply chain workflow) as well as a *Chief of Digital Robotics Officer* (to manage instead the automation of processes and activities).

Chapter 13
AI and Ethics

13.1 How to Design Machines with Ethically-Significant Behaviors

There has been a lot of talk over the past months about AI being our best or worst invention ever. The chance of robots taking over and the following catastrophic sci-fi scenario makes the ethical and purposeful design of machines and algorithms not simply important but necessary.

But the problems do not end here. Incorporating ethical principles into our technology development process should not just be a way to prevent human race extinction but also a way to understand how to use the power coming from that technology responsibly.

This chapter does not want to be a guide for ethics for AI or setting the guidelines for building ethical technologies. It is simply a stream of consciousness on questions and problems I have been thinking and asking myself, and hopefully, it will stimulate some discussion.

13.2 Data and Biases

The first problem everyone raises when speaking about ethics in AI is, of course, about data. Most of the data we produce (if we exclude the ones coming from observation of natural phenomena) are artificial creations of our minds and actions (e.g., stock prices, smartphone activity, etc.). As such, **data inherit the same biases we have as humans**.

First of all, what is a cognitive bias? The (maybe controversial) way I look at it is that a cognitive bias is a **shortcut of our brain that translates into behaviors which required less energy and thought to be implemented**. So, a bias is a good thing to me, at least in principle. The reason why it becomes a bad thing is that the

external environment and our internal capacity to think do not proceed *pari passu*. Our brain gets trapped into heuristics and shortcuts which could have resulted into competitive advantages 100 years ago but is not that plastic to quickly adapt to the change of the external environment (I am not talking about a single brain but rather on a species level).

In other words, *the systematic deviation from a standard of rationality or good judgment* (this is how bias is defined in psychology) is nothing more for me than a simple **evolutionary lag of our brain**.

Why all this *excursus*? Well, because I think that most of the biases data embed comes from our own cognitive biases (at least for data resulting from human and not natural activities). There is, of course, another block of biases which stems from pure statistical reasons (*the expected value is different from the true underlying estimated parameter*). Kris Hammond of Narrative Science merged those two views and identified at least five different biases in AI. In his words:

- **Data-driven bias** (bias that depends on the input data used);
- **Bias through interaction**;
- **Similarity bias** (it is simply the product of systems doing what they were designed to do);
- **Conflicting goals bias** (systems designed for very specific business purposes end up having biases that are real but completely unforeseen);
- **Emergent bias** (decisions made by systems aimed at personalization will end up creating bias "bubbles" around us).

But let's go back to the problem. How would you solve the biased data issue then?

Simple solution: you can try to remove any data that could bias your engine *ex-ante*. Great solution, it will require some effort at the beginning, but it might be feasible.

However, let's look at the problem from a different angle. I was educated as an economist, so allow me to start my argument with this statement: **let's assume we have the perfect dataset**. It is not only omni-comprehensive but also clean, consistent and deep both longitudinally and temporally speaking.

Even in this case, **we have no guarantee AI won't learn the same bias autonomously as we did**. In other words, removing biases by hand or by construction is not a guarantee of those biases to not come out again spontaneously.

This possibility also raises another (philosophical) question: we are building this argument from the assumption that *biases are bad* (mostly). So let's say the machines come up with a result we see as biased, and therefore we reset them and start again the analysis with new data. But the machines come up with a similarly 'biased result'. Would we then be open to accepting that as true and revision what we consider to be biased?

This is basically a **cultural and philosophical clash between two different species**.

In other words, I believe that two of the reasons why embedding ethics into machine designing is extremely hard is that (i) **we don't really know unanimously what ethics is**, and (ii) we should be open to admit that our values or ethics might not be completely right and that what we consider to be biased is not the exception but rather the norm.

Developing a (general) AI is making us think about those problems and **it will change** (if it hasn't already started) **our values system**. And perhaps, who knows, we will end up learning something from *machines' ethics* as well.

13.3   Accountability and Trust

Well, now you might think the previous one is a purely philosophical issue and that you probably shouldn't care about it. But the other side of the matter is about how much you **trust your algorithms**. Let me give you a different perspective to practically looking at this problem.

Let's assume you are a medical doctor and you use one of the many algorithms out there to help you diagnose a specific disease or to assist you in a patient treatment. In the 99.99% of the time the computer gets it right—and it never gets tired, it analyzed billions of records, it sees patterns that a human eye can't perceive, we all know this story, right? But what if in the remaining 0.01% of the case your instinct tells you something opposite to the machine result and you end up to be right? What if you decide to follow the advice the machine spit out instead of yours and the patient dies? Who is liable in this case?

But even worse: let's say in that case you follow your gut feeling (we know is not gut feeling though, but simply your ability to recognize at a glance something you know to be the right disease or treatment) and you save a patient. The following time (and patient), you have another conflict with the machine results but strong of the recent past experience (because of a *hot-hand fallacy* or an *overconfidence bias*) you think to be right again and decide to disregard what the artificial engine tells you. Then the patient dies. Who is liable now?

The question is quite delicate indeed and the scenarios in my head are:

(a) a scenario where the doctor is only human with no machine assistance. The payoff here is that liability stay with him, he gets it right 70% of the time, but the things are quite clear and sometimes he gets right something extremely hard (the lucky guy out of 10,000 patients);

(b) a scenario where a machine decides and gets it right 99.99% of the time. The negative side of it is an unfortunate patient out of 10,000 is going to die because of a machine error and the liability is not assigned to either the machine or the human;

(c) a scenario the doctor is assisted but has the final call to decide whether to follow the advice. The payoff here is completely randomized and not clear to me at all.

As a former economist, I have been trained to be heartless and reason in terms of expected values and big numbers (basically a **Utilitarian**), therefore scenario (b) looks the only possible to me because it saves the greatest number of people. But we all know is not that simple (and of course doesn't feel right for the unlucky guy of our example): think about the case, for instance, of autonomous vehicles that lose controls and need to decide if killing the driver or five random pedestrians (the famous ***Trolley Problem***). Based on that principles I'd save the pedestrians, right? But what about all those five are criminals and the driver is a pregnant woman? Does your judgement change in that case? And again, what if the vehicle could instantly use cameras and visual sensors to recognize pedestrians' faces, connect to a central database and match them with health records finding out that they all have some type of terminal disease?

The final doubt that remains is then not simply about liability (and the choice between pure outcomes and ways to achieve them) but rather on trusting the algorithm (and I know that for someone who studied 12 years to become doctor might not be that easy to give that up). In fact, **algorithm adversion** is becoming a real problem for algorithms-assisted tasks and it looks that people want to have an (even if incredibly small) degree of control over algorithms (Dietvorst et al. 2015, 2016).

But above all: **are we allowed to deviate from the advice we get from accurate algorithms**? And if so, in what circumstances and to what extent?

If an AI would decide on the matter, it will also probably go for scenario (b) but we as humans would like to find a compromise between those scenarios because we '*ethically*' don't feel any of those to be right. We can rephrase then this issue under the '**alignment problem**' lens, which means that the goals and behaviors an AI have need to be aligned with human values—an AI needs to think as a human in certain cases (but of course the question here is how do you discriminate? And what's the advantage of having an AI then? Let's therefore simply stick to the traditional human activities).

In this situation, the work done by the Future of Life Institute with **the Asilomar Principles** becomes extremely relevant.

The alignment problem, in fact, also known as '**King Midas problem**', arises from the idea that no matter how we tune our algorithms to achieve a specific objective, we are not able to specify and frame those objectives well enough to prevent the machines to pursue undesirable ways to reach them. Of course, a theoretically viable solution would be to let the machine maximizing for our true objective without setting it *ex-ante*, making therefore the algorithm itself free to observe us and understand what we really want (as a species and not as individuals, which might entail also the possibility of switching itself off if needed).

Sounds too good to be true? Well, maybe it is. I indeed totally agree with Nicholas Davis and Thomas Philbeck from WEF that in the Global Risks Report 2017 wrote:

> There are complications: humans are irrational, inconsistent, weak-willed, computationally limited and heterogeneous, all of which conspire to make learning about human values from human behaviour a difficult (and perhaps not totally desirable) enterprise.

What the previous section implicitly suggested is that not all AI applications are the same and that *error rates* apply differently to different industries. Under this assumption, it might be hard to draw a line and design an accountability framework that does not penalize applications with weak impact (e.g., a recommendation engine) and at the same time do not underestimate the impact of other applications (e.g., healthcare or AVs).

We might end up then designing **multiple accountability frameworks** to justify algorithmic decision-making and mitigate negative biases.

Certainly, the most straightforward solution to understand who owns the liability for a certain AI tool is thinking about the following threefold classification:

– *We should hold the AI system itself as responsible for any misbehaviour* (does it make any sense?);
– *We should hold the designers of the AI as responsible for the malfunctioning and bad outcome* (but it might be hard because usually AI teams might count hundreds of people and this preventative measure could discourage many from entering the field);
– *We should hold accountable the organization running the system* (to me it sounds the most reasonable between the three options, but I am not sure about the implications of it. And then what company should be liable in the AI value chain? The final provider? The company who built the system in the first place? The consulting business which recommended it?).

There is not an easy answer and much more is required to tackle this issue, but I believe a good starting point has been provided by Sorelle Friedler and Nicholas Diakopoulos. They suggest to consider accountability through the lens of five core principles:

– **Responsibility**: a person should be identified to deal with unexpected outcomes, not in terms of legal responsibility but rather as a single point of contact;
– **Explainability**: a decision process should be explainable not technically but rather in an accessible form to anyone;
– **Accuracy**: *garbage in, garbage out* is likely to be the most common reason for the lack of accuracy in a model. The data and error sources need then to be identified, logged, and benchmarked;
– **Auditability**: third parties should be able to probe and review the behavior of an algorithm;
– **Fairness**: algorithms should be evaluated for discriminatory effects.

13.4   AI Usage and the Control Problem

Everything we discussed so far was based on two implicit assumptions that we did not consider up to now: first, everyone is going to benefit from AI and everyone will be able and in the position to use it.

This might not be completely true though. Many of us will indirectly benefit from AI applications (e.g., in medicine, manufacturing, etc.) but we might live in the future in a world where only a handful of big companies drives the AI supply and offers fully functional AI services, which might not be affordable for everyone and above all not *super partes*.

AI democratization versus a centralized AI is a policy concern that we need to sort out today: if from one hand the former increases both the benefits and the rate of development but comes with all the risks associated with system collapse as well as malicious usages, the latter might be more safe but unbiased as well. Should AI be centralized or for everyone?

The second hypothesis, instead, is that we will be forced to use AI with no choice whatsoever. This is not a light problem and we would need a higher degree of education on what AI is and can do for us to not be misled by other humans. If you remember the healthcare example we described earlier, this could be also a way to partially solve some problem in the accountability sphere. If the algorithm and the doctor have a contradictory opinion, you should be able to choose who to trust (and accepting the consequences of that choice).

The two hypotheses above described lead us to another problem in the AI domain, which is the ***Control Problem***: if it is centralized, who will control an AI? And if not, how should it be regulated?

I wouldn't be comfortable at all to empower any government or existing public entity with such a power. I might be slightly more favourable to a big tech company, but even this solution comes with more problems than advantages. We might then need a new impartial organization to decide how and when using an AI, but history teaches us we are not that good in forming mega impartial institutional players, especially when the stake is so high.

Regarding the AI decentralization instead, the regulation should be strict enough to deal with cases such as **AI-to-AI conflicts** (what happens when 2 AIs made by two different players conflict and give different outcomes?) or the ethical use of a certain tool (a few companies are starting their own **AI ethics board**) but not so strict to prevent research and development or full access to everyone.

I will conclude this section with a final question: I strongly believe there should be a sort of '***red button***' to switch off our algorithms if we realize we cannot control it anymore. However, the question is who would you grant this power to?

13.5   AI Safety and Catastrophic Risks

As soon as AI will become a commodity, it will be used maliciously as well. This is a virtual certainty. And the value alignment problem showed us that we might get in trouble due to a variety of different reasons: it might be because of misuses (**misuse risks**), because of some accident (**accident risks**), or it could be due to **other risks**.

But above all, no matter the risk we face, it looks that AI is dominated by some sort of exponential chaotic underlying structure and getting wrong even minor things could turn into catastrophic consequences. This is why is paramount to understand every minor nuance and solve them all without underestimating any potential risk.

Amodei et al. (2016) actually dug more into that and drafted a set of five different core problems in AI safety:

– **Avoiding negative side effects**;
– **Avoiding reward hacking**;
– **Scalable oversight** (respecting aspects of the objective that are too expensive to be frequently evaluated during training);
– **Safe exploration** (learning new strategies in a non-risky way);
– **Robustness to distributional shift** (can the machine adapt itself to different environments?).

This is a good categorization of AI risks but I'd like to add the *interaction risk* as fundamental as well, i.e., the way in which we interact with the machines. This relationship could be beneficial (see the *Paradigm 37–78*) but comes with several risks as well, as for instance the so-called *dependence threat*, which is a highly visceral dependence of human on smart machines.

A final food for thought: we are all advocating for full transparency of methods, data and algorithms used in the decision-making process. I would also invite you though to think that full transparency comes with the great **risk of higher manipulation**. I am not simply referring to cyber-attacks or bad-intentioned activities, but more generally to the idea that once the rules of the game are clear and the processes reproducible, it is easier for anyone to hack the game itself.

Maybe companies will have specific departments in charge of influencing their own or their competitors' algorithms, or there will exist companies with the only scope of altering data and final results. Just think about that…

13.6   Research Groups on AI Ethics and Safety

There are plenty of research groups and initiatives both in academia and in the industry start thinking about the relevance of ethics and safety in AI. The most known ones are the following 20, in case you like to have a look at what they are doing:

(a) Future of Life Institute (Boston);
(b) Berkman Klein Center (Boston);
(c) Institute Electrical and Electronic Engineers—IEEE (Global);
(d) Centre for the study on existential risks (Cambridge, UK);
(e) K&L gates endowment for ethics (Global);
(f) Center for human-compatible AI (Berkeley, CA);
(g) Machine Intelligence Research Institute (Berkeley, CA);
(h) USC center for AI in society (Los Angeles);
(i) Leverhulme center for future of intelligence (Cambridge, UK);
(j) Partnership on AI (Global);
(k) Future of Humanity Institute (Oxford, UK);
(l) AI Austin (Austin, US);
(m) Open AI (San Francisco, US);
(n) Campaign to Stop Killer Robots (Global);
(o) Campaign against Sex Robots (Global);
(p) Foundation for Responsible Robotics (Global);
(q) Data & Society (New York, US);
(r) World Economic Forum's Council on the Future of AI and Robotics(Global);
(s) AI Now Initiative (New York, US);
(t) AI100 (Stanford, CA).

Finally, Google has just announced the 'People + AI research' (PAIR) initiative, which aims to advance the research and design of people-centric AI systems.

13.7 Conclusion

Absurd as it might seem, I believe ethics is a technical problem. Writing this post, I realized how much little I know and even understand about those topics. It is incredibly hard to have a clear view and approach on ethics in general, let's not even think about the intersection of AI and technology. I didn't even touch upon other questions that should keep AI experts up at night (e.g., unemployment, security, inequality, universal basic income, robot rights, social implications, etc.) but I will do in future posts (any feedback would be appreciated in the meantime).

I hope your brain is melting down as mine in this moment, but I hope some of the above arguments stimulated some thinking or ideas regarding new solutions to old problems.

I am not concerned about robots taking over or Skynet terminates us all, but rather of humans using improperly technologies and tools they don't understand. I think that the sooner we clear up our mind around those subjects, the better it would be.

References

Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). *Concrete Problems in AI Safety*. arXiv:1606.06565v2.

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology, 144*(1), 114–126.

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2016). *Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them*. Available at SSRN: https://ssrn.com/abstract=2616787 or http://dx.doi.org/10.2139/ssrn.2616787.

Chapter 14
AI and Intellectual Property

14.1 Why Startups Patent Inventions (and Why Is Different for AI)

Working alongside early-stage companies, venture funds, and corporations, I often wonder what the word "**defensibility**" means nowadays for an AI startup. And even if I find myself often thinking and advising people on why a patent on a machine learning product/algorithm makes sense or not, I recently realized that this problem is actually related to how I see an AI company be protected in the long-run.

Patenting an invention is one of the four forms of IP protection (i.e., patents, trademarks, copyrights, and trade secrets) and it basically gives you the chance to use and exploit the economic benefits of a certain invention for a quite long period (usually 20 years) demanding in turn that you make your invention public for the sake of the scientific and technological progress.

For a startup, having a patent has historically been a great advantage over the competition, but is also a burdensome cost to cover especially at a very early stage (around 20k both in Europe and US over a 3–5 years period). And this cost is not even getting into account the fact that obtaining a patent is the same of tossing a coin (only 55% of the applications result into a granted patent, eventually—see Carley et al. 2014 for more details).

So why companies get themselves into this tricky and expensive process in the first place? Well, there are multiple reasons for doing it. You can indeed create a strong competitive advantage or shape a new stream of revenues by licensing your technology, or simply increase your confidence that what you are building is being recognized by the rest of the world as useful.

Indeed one of the companies I am working with called Meeshkan ML recently applied for a patent for their distributed machine learning algorithm because in the founder's words "*we feel that we can share it with our local community of engineers and build new algorithms on top of it without worrying about losing our*

business advantage. We therefore made a calculated risk in investing time into the R&D necessary for the patent instead of pushing a product out."

But most of all, a startup often heavily depends on external financing—and investors love(d) patents. It is a simple way for many of them to achieve three things at once: they get sure the tech is legit (i.e., ***they outsource the technical due diligence*** to patent lawyers and offices); they are more confident the ***technology is feasible*** (reducing product risks); they get more confident the ***team can actually build*** what they claimed and that is committed to it (reducing execution and team risks).

This is especially true for deep science ventures as well as emerging technologies where understanding those three aspects is incredibly cumbersome. AI is clearly one of those fields, but I am also claiming it stands apart from other technological inventions for a bunch of different reasons:

- **A strong open-source community**: many AI algorithms/libraries/packages are completely open-sourced, and you usually tend to build on top of those (which makes patenting extremely hard to be managed);
- **Confusion around the (real) invention**: it is still hard to identify which part of an AI algorithm is the real source of invention. Is it the source code? Or rather the data used? Or maybe the process or the human-machine blending interaction?
- **Continuous evolution**: feedback loops in machine learning push the code to keep evolving, which complicates the understanding of whether a new invention breaks the initial patent.
- **It stems from an academic community**: traditionally many inventions are industry-driven while AI is historically intrinsically related to academia, where the approach to scientific research and development is culturally very different;
- **It is still half science half art**: even though a patent requires a fairly decent degree of details around the invention in question, when it comes to AI the devil is into the smallest details (e.g., tuning). In other words, even with a full access to the process or to the algorithm, you might not be able to use it and implement it correctly.

14.2   The Advantages of Patenting Your Product

So, if you are an AI entrepreneur, you might now question the utility of a patent and above all interrogating yourself with the big question: **should I even try to get a patent in the first place** or invest my money in more meaningful activities?

Quick answer: **it depends**.

In addition to common sense reasoning (which is more or less what I mentioned above), there is empirical evidence that a patent may be beneficial to your business.

In fact, Mann and Sager (2007) proved that **having patents is positively correlated to startup performance indicators**, such as the number of financing rounds, total investment received, exit status, longevity and late-stage financing. This is not an isolated example since Conti et al. (2013) for instance also showed that **having a patent increases the likelihood of getting VC funding** (but no angel funding, because angels apparently value other things more than IP protection). Haeussler et al. (2014) interpreted patents as a quality signals for VC financing, finding that **one patent application reduces time to venture funding by 76%** (and VCs are apparently very good at understanding the quality of a patent) and therefore increases the likelihood of getting funding (an interesting side evidence is that also patent oppositions are well seen and increase likelihood of funding, maybe because they are signs of market potential and validation). This probability is even higher when you complement patents with trademarks (Zhou et al. 2016).

Furthermore, Baum and Silverman (2004) also found a **positive association between patent applications and pre-IPO financing**, while Farre-Mensa et al. (2017) instead pushed this analysis one step forward. They did not simply show that the first patent you obtain **increases your likelihood of raising venture capital funding in the following three years by 2.3** percentage points (a 53% increase relative to the 4.3% unconditional probability of raising VC funding, although the effect is smaller during the first year and ramps up later) as well as the likelihood of securing a loan by pledging the patent as collateral, but also showed the impact of them on the daily operations.

In fact, startups that file a patent have, on average, **55% higher employment growth** and **80% higher sales growth five years later**. They also pursue more **follow-on quality innovation** and have a higher probability of getting subsequent patents (up by 49%) of a higher quality (with the average number of citations per subsequent patent increasing by 26%).

Eventually, have a portfolio of patents will have a strong positive effect also in case of an IPO.

> Hence, let's quickly recap what we learned: a patent may bring you more funding, more quick funding, attract talents and increase sales, put you on the right road to keep building great things and increase your exit value.

Well, truth be told, it sounds like a great benefits package to me. So why is not everyone rushing to patent every single invention developed? Here comes the catch.

First of all, the results that we mentioned above do not apply to everyone (Farre-Mensa et al. 2017). Generally speaking, they apply only to startups that **have raised little or no money before patenting** (very early-stage), lead by **inexperienced entrepreneurs** (serial entrepreneurs do not need patent to prove the company potential and tech feasibility), mostly **operating in IT**, and located in **hot technology hubs** (where getting investors' attention is otherwise difficult).

Moreover, **additional patenting through funding rounds do not increase the investment amounts**.

Block et al. (2013) indeed show that an **inverted U-shaped relationship** exists between the number of trademarks and the company's financial valuations

(i.e., it is positively related up to a certain threshold and then becomes negative), and only for early-stage startups. The same type of relationship also exists between the breadth of start-ups' trademark portfolios and the financial valuations, and it is also valid for the relationship between diversified patenting activity and performance (Fernhaber and Patel 2012; Li et al. 2012; Qian 2002).

> So, if you are an early-stage startup trying to find your place in the world and equipped with a brilliant technical innovation in your hand, I would at least encourage you to have this conversation with a patent lawyer. The first patent you filed is definitely worth its value.

14.3 Reasons Behind not Looking for Patent Protection

Of course is not all peaches and dandelions here. If from one hand patents increase valuation and attract investors' attention, some studies seem to suggest otherwise (Smith and Cordina 2014). Looking at the patents portfolio is a shortcut for less sophisticated investors, but for the best one is nothing more than an extra data point in their complex risk-return simulation modeling.

Many startups deliberatively choose then to not patent any internally generated innovation, and most of the time they do so for one or more of the following reasons:

- **They are too early-stage**: the majority of companies working nowadays in AI are relatively young (according to CBinsights, 69% or more of AI deals since 2012 have gone to startups still in the early stages). This often means they have no means, either financial or in terms of resources and time, to file a patent;
- **They might convey too much information**: even though a single patent might not say much about your business, a portfolio of them can actually give many more strategic insights to a careful observer (and probably more than what you aim for);
- **The patent landscape is messy**: as we already mentioned above, the presence of widespread open-source libraries drastically increases the complexity of filing a patent;
- **Patents are becoming less meaningful**: it is extremely easy to violate a patent without even knowing it, and incredibly hard to enforce one. Add to this problem the corollary issue that patent litigation (when you are actually able to recognize a violation, which is not obvious) costs you around half a million (a trademark instead between 300 and 500 k) and you will not want to get yourself into that problematic spot at all;
- **Cultural divergence** (i.e., they do not believe in it): I might spend many words here, but I rather prefer to quote Tesla announcement of a couple of years ago:

> Tesla will not initiate patent lawsuits against anyone who, in good faith, wants to use our technology. […] Technology leadership is not defined by patents, which history has repeatedly shown to be small protection indeed against a determined competitor, but rather by the ability of a company to attract and motivate the world's most talented engineers. We believe that applying the open source philosophy to our patents will strengthen rather than diminish Tesla's position in this regard (Elon Musk).

- **They have alternative protections**: you can easily supply the absence of patents with different traditional and non-traditional moats (by the way, if you haven't already, go and read Gil Dibner's post and Jerry Chen's one of this topic):

  (i) *Data moat*: most of the AI software need to be fed with millions of data points. Even if you have access to the algorithm, without data you cannot basically use it. Furthermore, having access to different datasets (or multiple "*systems of records*", as Jerry Chen called them) and use them together not only provides you with an exponentially higher value but also increases the likelihood you can deeply personalize your product and tailor it to your customers (which then increases their **switching costs** toward the competition);

  (ii) *Network effect* (also called *Metcalf's law* or *AI flywheels*): the value of the network increases exponentially with the number of the nodes, which in other words means that new customers imply more data, which implies a better algorithm, which in turn implies new customers, and so on so forth. You need though to reach a *minimum algorithmic performance* (MAP) before the network can attract the first cohort of customers with a strong value proposition;

  (iii) *Talents*: AI as a field is dominated by academic-trained talents, which are incredibly expensive and hard to hire and often demands an open approach to research (i.e., open-source software and publications in scientific journals). There are very few strong AI researchers worldwide, and securing one major name and/or a good team can represent a strong barrier for a startup (although you really hit the jackpot if you hire someone with specific **domain expertise**, which is quickly becoming the real differentiation between succeeding and surviving);

  (iv) *Cost efficiency*: many of the most powerful AI-driven businesses today are able to scale while limiting costs related to teams and product development (and let's not forget that until some time ago the median team size for startups acquired by big incumbents was around 7).

Again, patenting is not a solution that works for everyone, but if you have none of the moats above I would highly recommend you to speak to a patent lawyer and look for some additional IP protection.

14.4   The Patent Landscape

If you are an AI entrepreneur that wants to try to patent something, it is certainly useful to start from mapping the landscape of existing patents to see whether something similar has already been filed and whether an opportunity materializes for your innovation.

This is a very specific process related to what you want to build and goes beyond the scope of this post, but what it can be intriguing to present is the current AI spectrum of patents and players.

CB Insights analyzed over 1150 AI companies in the last decade (since 2009) finding that **21% applied** for a patent and **only 11% eventually got** at least one. In particular, in the last five years, the issue is becoming more tricky and the US IP law started drawing some clear lines (the most famous case up to date is *Alice Corp. v. CLS Bank*, where they rejected a patent application on computer software because too *abstract*). Nowadays, it seems therefore clear that training sets, proprietary information, a particular expression of source code, and many other steps in an AI value chain *cannot be patented*. On the other side though, *trade secret protection* is suitable for a wide variety of circumstances (e.g., neural networks, training sets, AI-generated code, learning algorithm, etc.).

CBinsights further divided the **patents of the "Big 5"** (Apple, Google, Amazon, Facebook, and Microsoft) from the **patents obtained by startups.**In this fashion, big tech giants patented several things, although many times because the **software is often attached to a hardware component** (Amazon for logistic robotics, Apple for the iPhone, and even Google for their smartphone products). Google leads the pack in applying for AI patents, while Microsoft is the most prolific in filing for patents overall. We are talking here about a relatively small magnitude (less than 40 AI patents in 2017, with a peak of 164 in 2015). On the startup side instead, Cortica and Numenta are dominating this space (with 38 and 37 patents, respectively) followed by Butterfly Network (27), SoundHound (26) and Smart Drive (24). Interestingly enough, in addition to traditionally strong-IP fields (e.g., healthcare), most of the patents have been filed on **horizontal AI type of applications** (in other words, platform, neuroscience approach to AI, GAI, core AI, etc.).

If we then widen the spectrum of what is being patented worldwide at any level and from any organization, it seems clear that most of the AI patents focus on **enabling intelligent robots** (e.g., self-driving cars, automated delivery drones, AI assistants, etc.), **deep learning**, **face recognition** and **AI hardware** (especially in China).

Finally, US and China are the two countries where most of this innovation (over 50% of the patents registered) has been happening over the last decade or so (which one of the two is leading is still controversial though), followed by Japan, South Korea, Germany, Canada, the U.K., Australia, India, and Russia. China seems to have an uncontested supremacy over *deep learning and machine vision*, where US is more keen to develop *NLP and other machine learning* technologies, and it is filing up patents at a much faster pace than its American counterparts (interestingly

enough, Chinese researchers more than doubled the number of AI scientific papers published in 2017 compared with 2010, trend that is going downward in the US). Europe instead places itself in the middle, with a flattening growth in patents applications (OECD 2017) and accounting for 10% of the US quantity. However, even if patenting is not so widespread in Europe, the scientific publication side is steadily growing (and of course, many recent announcements have been done to better position Europe in the AI race).

14.5   Conclusions

When I started thinking about this IP issue, my main question was "how is it possible to patent something when the market is driven by open-source forces?", cause I was naively assuming the two elements being incompatible. The reality is that they are not because they might serve different purposes and appear in different stages of a company life.

Bottom line, keep your mind open about the chance of protecting your IP when you are very early-stage or to send a signal out there about what you are building, but do not waste too much money and time on it if you are already building a good traction quickly. Do not also assume that a patent will vouch for everything you say and do because it won't, and it is not a substitute to a great co-founding team or a strong market validation, and by all means not having a patent does not imply your product is inferior or your company is worse than others.

To say that with terms that we all know, "***patents are vitamins, not painkillers***".

References

Baum, J. A., & Silverman, B. S. (2004). Picking winners or building them? Alliance, intellectual, and human capital as selection criteria in venture financing and performance of biotechnology start-ups. *Journal of Business Venturing, 19,* 411–436.

Block, J. H., De Vries, G., Schumann, J. H., & Sandner, P. (2013). Trademarks and venture capital valuation. *Journal of Business Venturing, 29*(4), 525–542.

Carley, M., Hegde, D., & Marco, A. (2014). What is the probability of receiving a U.S. patent? *Yale Journal of Law and Technology, 17*(1), 204–224.

Conti, A., Thursby, M., & Rothaermel, F. T. (2013). Show me the right stuff: Signals for high-tech startups. *Journal of Economics & Management Strategy, 22*(2), 341–364.

Farre-Mensa, J., Hedge, D., & Ljungqvist, A. (2017). *What is a Patent Worth? Evidence from the US Patent "Lottery".* NBER Working Paper, 23268: 1–57.

Fernhaber, S. A., & Patel, P. (2012). How do young firms manage product portfolio complexity? The role of absorptive capacity and ambidexterity. *Strategic Management Journal, 33*(13), 1516–1539.

Haeussler, C., Harhoff, D., & Mueller, E. (2014). How patenting informs VC investors—The case of biotechnology. *Research Policy, 43*(8), 1286–1298.

Li, L., Qian, G., & Qian, Z. (2012). The performance of small and medium-sized technology-based enterprises: Do product diversity and international diversity matter? *International Business Review, 21*(5), 941–956.

Mann, R. J., & Sager, T. W. (2007). Patents, venture capital, and software start-ups. *Research Policy, 36,* 193–208.

OECD. (2017). Science, innovation and the digital revolution. In *OECD science, technology and industry scoreboard 2017: The digital transformation*. Paris: OECD Publishing.

Qian, G. (2002). Multinationality, product diversification, and profitability of emerging US small- and medium-sized enterprises. *Journal of Business Venturing, 17*(6), 611–633.

Smith, J. A., & Cordina, R. (2014). Patenting and the early-stage high-technology investor: Evidence from the field. *R&D Management, 45*(5), 589–605.

Zhou, H., Sandner, P. G., Martinelli, S. L., & Block, J. H. (2016). Patents, trademarks, and their complementarity in venture capital funding. *Technovation, 47,* 14–22.

Chapter 15
AI and Venture Capital

15.1 The Rationale

Great startups very often stem from a simple idea and a couple of people that identify a specific need and a solution to address an existing market gap. On the other side of the fence, there are angels and venture capitalists (VC), who provide the capital and sometimes the necessary help to make that dream become a reality.

If you are an investor then, there are two things you want to know all the time: **where the good companies are** (*scouting abilities*) and how to know **whether a specific company will be a good investment** or not (*cherry-picking ability*). The scouting abilities are usually refined through years of networking and branding efforts, while cherry-picking skills are often emerging from pattern-matching, i.e., seeing company over company year over year and trying to intuitively infer and deduce why a startup succeeds while others do not.

As it is clear, the VC investment process is incredibly slow, labor-intensive, inefficient, expensive, and often even biased (Sheehan and Sheehan 2017). Many venture capitalists suffer indeed from common psychological biased such as **overconfidence** (Zacharakis and Shepherd 2001); **availability biases** (over-weighting information that comes easily to mind because memorable while underweighting information that is less exciting); **information overload** (Zacharakis and Meyer 2000), meaning that more information often leads only to greater confidence and not to greater accuracy; **halo effect** (*how similar this company is to previous exits I had?*); **survivorship biases**; **representativeness**, which means ignoring statistical information in favor of a narrative; **confirmation bias** (accepting information that support pre-existing beliefs); and **similarity biases** (meaning not simply that entrepreneurs with similar educational and professional path are preferred, but also that VCs with a history of working with startups tend to overlook the potential of entrepreneurs with a background in established firms, and vice versa—Franke et al. 2006).

Given then the incredible degree of uncertainty that this business embeds, it might be worth to look for some external help—let's say a more "*automatic*" one. However, **men and machines are equally bad at predicting success outcomes** (Mckenzie and Sansone 2017), and although I am not even sure that this effort is a feasible one, I am convinced that something valuable can be found in previous (partial) attempts to understand the basics of a successful business.

15.2   Previous Studies

The literature on venture capital is quite vast, and it covers a variety of sub-topics ranging from investment choices and exits to organizational issues, relationships, contracting, post-investment, and much more (Da Rin et al. 2013).

My goal here is only to focus on studies that investigate the impact of certain variables on the likelihood of success of an early-stage company, either direct or indirect. In order to do it, I am grouping different studies in clusters that represent the source of a specific competitive advantage that increases the likelihood of an exit: **personal and team characteristics**, **financial considerations**, and **business features**.

15.2.1   *Personal and Team Characteristics*

This group concerns all the traits that are strictly related to the entrepreneur or to the founding team. Starting with mere demographics, McKenzie and Sansone (2017) show that **people in their 30s**—and in particular male that scored highly on ability tests—are more likely to succeed (findings confirmed also in McKenzie and Paffhausen 2017). Moreover, being previously **unemployed** and **being married/cohabit** are respectively negatively and positively correlated with a higher success rate (Miettinen and Littunen 2013).

Previously Gompers et al. (2010) showed that **being a successful serial entrepreneur** increases the success rate in future new ventures, although this is no longer true for entrepreneurs that were previously unsuccessful. In particular, a good serial entrepreneur seems to be better than average at picking the right industry and the time to start a new company (i.e., timing the market). Hsu (2007) does not only support the findings of Gompers et al. (2010) but also provides the evidence that **serial entrepreneurs get better valuations**.

> *A lower valuation from a highly reputable VC is on average preferred to obtain a higher valuation from a low-reputation investor.*

However, it is not only the experience that matters in this industry. Bengtsson and Hsu (2010) indeed show that **ethnic similarity**, as well as **attendance**

of the same top universities, increase the likelihood of receiving funding (findings supported also by Sunesson 2009), while Shane and Stuart (2002) prove that **_social capital_** (i.e., having direct or indirect ties with reputable venture capitalists) improves your chances when fundraising. Hsu (2004) even shows that getting a _lower valuation from a highly reputable VC is on average preferred to obtain a higher valuation from a low-reputation investor._

Social networks are also very relevant, whether online or offline. In this fashion, Nann et al. (2010) find that the success of a company is intrinsically connected to its founder's network robustness. Hence, if a founder comes out from a top university and consistently maintain her links with alumni from that same university, she is more likely to become successful. Gloor et al. (2011, 2013) instead analyzed the entrepreneurs' email traffic and social media activity to understand whether this correlates with the startup success and eventually found that **_centrality_** to the network increases the probability of an exit rate. Finally, whether those networks encompass only strictly personal relationships or external business and more formal ones, it does have an impact depending on the stage of the company (respectively, in the first four years external networks have a positive impact on company performance while the same is true in the following four years for internal networks—Littunen and Niitykangas 2010).

Eesley et al. (2014) focus instead their attention on team composition finding that a **diverse team exhibits a higher performance**. However, this does not happen all the time but only when in a competitive commercialization environment. On the other hand, technically focused founding teams are more effective when in a cooperative commercialization environment and when pursuing an innovation strategy. However, on the other hand, Mueller and Murmann (2016) investigated the complementarity of skills in the human capital base of a startup (i.e., co-founding team and employees) finding that the mix of business and technical skills has an exponential impact on the company performance only when the founder has technical knowledge and employs additional business experts (not the other way round, and neither when business and technical skills are balanced within a founding team).

Sometimes though, the founder is not the right person to keep running the company after a certain stage. Ewens and Marx (2017) indeed prove that **replacing the founder/s with experienced managers can often improve the company performance**.

Finally, there are traits that slightly more "_esoteric_", in the sense that is quite hard to understand whether a non-spurious correlation exists between a certain factor and the company performance. Entrepreneurship literature has indeed given attention even to the signal originated from calling a company after the owner name. If this results into a more successful company for reputational cost reasons (Belenzon et al. 2017) or into a non-performing firm because of its not ambitious growth-oriented mindset (Guzman and Stern 2014) is still to be decided.

15.2.2 Financial Considerations

This group embeds every research I could find on the impact of funding and financial variables on predicting the success of a company. From a financial point of view, Miettinen and Littunen (2013) studied the impact of **equity share** (i.e., the capital owned by the entrepreneur as a percentage of total assets) finding that it has some predictive power over the probability that the company will do better than its competitors.

On the funding side instead, Puri and Zarutskie (2012) showed that **VC-backed companies are more likely to go public or be acquired** and less likely to fail, which is also supported by the findings of Hsu (2006), Nahata (2008), Sorensen (2007) and Inderst and Mueller (2009).

Cumming (2008) shows instead that the probability of **being acquired is positively correlated with having obtained financing through convertible securities** rather than equity deals, while the opposite is true for IPOs.

Zarutskie (2010) also proved that if the **partners of a venture capital fund have prior experience either in VC or startups**, they are more likely to outperform their competitors (in terms of portfolio companies exited), and the same is also true for partners with prior industry experience (and, interestingly enough, **having an MBAs is somehow negatively correlated with the fund performance**). Gompers et al. (2009) also found that **VC firms that are specialized in a few industries perform better** than generalist VCs, while Ewens and Rhodes-Kropf (2015) showed that there exists a sort of **persistence in the performance of a partner of a VC firm**. If the partner has brought companies to an IPO in the past, she will keep doing it with future startups. If she instead preferred the acquisition path, more acquisitions will come later for her own portfolio companies—and if she failed multiple times, she will keep failing.

Tian (2011) found instead that **syndicate deals are more likely to produce an exit** and to do it at a higher valuation. Miloud et al. (2012) also showed that a higher valuation can be reached through a **higher product differentiation**, **industry growth rate**, **completeness of the management team**, and based on whether the **founders had previous industry** and **management experience**.

Finally, it is important to consider the investment critical mass, if any—i.e., a certain threshold of funding that increases the likelihood of being successful. Lasch et al. (2007) and Groenewegen and de Langen (2012) showed indeed that **raising more than €75,000 results in a greater chance to outperform** your peers.

15.2.3 Business Features

This group includes instead other factors that explain the differential performance of a company and that are related to intrinsic characteristics of the business (e.g., technology, intellectual property, etc.).

Different studies (Cockburn and MacGarvie 2009; Mann and Sager 2007; Hsu and Ziedonis 2011) prove empirically that the higher the **number of patents** a company has, the more likely it is to obtain venture financing and to exit through either an IPO or an acquisition. This is especially true for early-stage companies, where the financial information is either missing or simply forecasted.

Lindsey (2008) showed that **strategic alliances are associated with a higher exit rate**, and similar results are presented also in Hoenig and Henkel (2015).

Halabi and Lussier (2014) found that having **clear financial and accounting information**, as well as a certain degree of entrepreneurial attitude and an **adequate working capital** are positively correlated with a higher likelihood of success, while in a later study Marom and Lussier (2014) proved this likelihood to be positively associated with having **professional advice**. The study is then run across different countries (Lussier and Halabi 2010) and extended to 26 independent variables in Teng et al. (2011).

15.2.4   Industry Knowledge

All the information presented so far comes directly from academic studies or research-oriented projects. There are tough important signals that the industry has capture over the years which are worth to be mentioned.

In particular, a research made by First Round some time ago shows that having at least one female co-founder increases the probability of success and performance of a company.

David Coats from Correlation Ventures instead recently released an analysis where he shows that having more than two VCs on the board is counter-productive, as well as having none.

Although those are not peer-reviewed results, given the data-driven approach and the brand of the two firms, I found them quite plausible and I am then including them in this list.

Hence, to recap what learnt so far, everything has been summarized in Table 15.1.

15.2.5   An Outsider Study: Hobos and Highfliers

*Even though is clearly a competitive advantage for a VC to know whether a company will be successful or not, this does not guarantee that the entrepreneur will want your money. **Timing the market** and **establishing a strong relationship with a founder** earlier on is as much important as knowing whether a company is more likely to succeed than not.*

In this light, one study caught my eyes (and very likely it is a unique research of its kind), which was a joint study between Haas Business School and Bloomberg

Table 15.1 Taxonomy of signals to predict probability of startup success

Impact	Variable	Reference	Effect
Direct	Track record	Gompers et al. (2010); Hsu (2007)	Previous successful entrepreneurs are more likely to succeed
	Patents	Cockburn and MacGarvie (2009); Mann and Sager (2007); Hsu and Ziedonis (2011)	The higher the number of patents, the higher the likelihood of exit
	Investors' Brand	Hsu (2006); Nahata (2008)	If financed by more reputable VCs, increases likelihood to do exits
	Convertible Notes	Cumming (2008)	Acquisition more likely (and IPO less likely) if used convertible notes
	Strategic Alliances	Lindsey (2008); Hoenig and Henkel (2015)	Strategic alliances associated with higher exit
	School	Sorenson (2009)	Exit likelihood increases if VC and entrepreneur went to top 3 schools
	Syndicate	Tian (2011)	Syndicate deals are more likely to exit at higher value
	Professional Advice	Marom and Lussier (2014)	Better likelihood if have professional advisors
	Team Diversity	Easley et al. (2013; 2014)	Diverse team exhibits higher performance if in competitive markets
	Network	Gloor et al. (2013; 2016); Nann et al. (2010)	If entrepreneur has a strong network and is central with respect to that network, higher success
	Working capital	Halabi and Lussier (2014)	If company has adequate working capital, more likely to succeed
	Accounting information	Halabi and Lussier (2014)	If company has clear financial/accounting statements, more likely to succeed
	Parents	Halabi and Lussier (2014)	If entrepreneur has parents who were entrepreneurs, higher success
	VC Support	Puri and Zarutskie (2012); Sorensen (2007) ; Inderst and Mueller (2009)	VC-backed companies are more likely to go public or acquired and less likely to fail.
	Founders age	McKenzie and Sansone (2017); McKenzie and Paffhausen (2017)	Founders in their 30s are more likely to succeed (alternatively, 8 years of experience post-college)
	Marital Status	Miettinen and Littunen (2013)	If founder is married or cohabit, positively correlated with likelihood of success
	Previous Employment	Miettinen and Littunen (2013)	If founder is coming from being previously unemployed, negative correlation with future success
	Founders Replacement	Ewens and Marx (2017)	If the founders are replaced with experienced managers, performance improves
	Equity Share	Miettinen and Littunen (2013)	The equity share, i.e. the capital owned by the entrepreneur as a percentage of total assets, is positively associated with likelihood of success
	VC Partner Experience	Zarutskie (2010)	If partners of a VC fund have prior experience either in VC or startups, they have greater and more exits
	VC Partner Specialization	Gompers et al. (2009)	VC firms that are specialized in a few industries have more exits than generalist VCs
	VC Partner Exits	Ewens and Rhodes-Kropf (2015)	If partners has done IPOs in the past, likelihood to IPO increases. The same is true for acquisition
	Marketing	Halabi and Lussier (2014)	If company has strong marketing efforts, more likely to succeed
	Minimum Investment	Lasch et al. (2007); Groenewegen and de Langen (2012)	Having raised more than €75,000 increases the probability of success
Indirect (deal formation)	Social Capital	Shane and Stuart (2002)	If entrepreneur has direct or indirect ties with reputable VC, can get funding easier
	Ethnicity/School	Bengtsson and Hsu (2010)	Similar ethnicity startup VC or university increase likelihood of investment
	Reputation	Hsu (2004)	High-reputation-low-valuation VCs are preferred to low-reputation-high-valuation investors
Indirect (valuation)	Product Differentiation	Miloud et al. (2012)	The higher the product differentiation, the higher the valuation
	Industry Growth	Miloud et al. (2012)	The higher the industry growth rate, the higher the valuation
	Industry Experience	Miloud et al. (2012)	If founders have industry experience, higher valuation
	Management Experience	Miloud et al. (2012)	If founders have previous management experience, higher valuation
	Management Team	Miloud et al. (2012)	If management team complete, higher valuation
	Management Background	Mueller and Murmann (2016)	Company performs better if management team is technical and employs business people, not the other way round
Industry knowledge	Gender	First Round Capital	Companies with at least 1 female founder perform better
	Board composition	Correlation Ventures	If you have more than 2 VCs on the board, exits will be lower (controlled for investment stages, industry groups, and time periods). The same is true for board with no VC

Beta (a Silicon Valley VC). **The researchers tried to predict ex-ante who will start an interesting venture and reshaped the idea we have on founders' stereotypes** (Ng and Stuart 2016).

In fact, potential future founders **have at least 8 years of experience after college** and start a company mainly during **market booms**. Furthermore, the **higher the education degree**, the lower probability to transition to self-employment but the higher the probability to become an entrepreneur (having a Master's degree, or a Ph.D. in a lighter way, represents the average for US founders). Even though a **technology background is often preferred, is not strictly required** (an MBA drastically increases the likelihood of starting a company, for instance). Finally, if you hold a position that spans both **technical and managerial responsibilities** it is more likely that you will want to start a new venture someday soon.

15.3   Who Is Using AI in the Private Investment Field

Assuming all the material presented above is groundbreaking would be short-sighted and thinking to be the first one to propose a data-driven approach to investing would be, if not only pretentious, also completely false.

In this respect, there are a bunch of funds and individuals out there who are actively looking at the same space. Clearly, knowing exactly what they do is quite impossible without having inside information, so I am simply reporting what I read and heard:

- **Correlation Ventures**: probably the first real data-driven investor, it reaches a decision on whether to invest or not in 2 weeks, plus other 2 for extra due diligence. They only do co-investments in the US and don't take board seats;
- **EQT Ventures**: with more than €500 M AuM and equipped with an AI system called "the Motherbrain", they have done more than 20 investments in less than two years;
- **SignalFire**: the firm run by Chris Farmer does not only use analytics to pick the right companies but also to help them grow by providing market intelligence and talent matching services;
- **WR Hambrecht Ventures**: Thomas Thurston is the key man behind WRH (and Growth Science, its sister tech company) who is advocating for the use of data science to guide growth investments;
- **venture/science**: a quant-driven VC led by Matt Oguz, it uses AI and decision theory to compute the risk associated with different attributes such as team completeness, vision, etc.;
- **645 Ventures**: Series A investor, it follows a metrics-driven approach to Growth Seed investing in a bunch of different sectors;
- **Hone Capital**: the Palo Alto-based US arm of CSC Group, it partnered with AngelList to create their proprietary model;
- **InReach Ventures**: led by Roberto Bonanzinga, InReach has quickly built a name as the software powered house able to scout early-stage European startups even before others VCs realized they need funding (see what happened with Oberlo, for example). Apparently, it took them two years and £5 m in investment to build their proprietary software;
- **Fly Ventures**: the recently closed a first €41 M fund to do small ticket size investments (up to €1 M) and have invested in companies like Bloomsbury AI, recently acquired by Facebook;
- **Social Capital**: led by Chamath Palihapitiya, the firm is better known to have started the *Capital as a Service (CaaS)* concept, and more recently they created an analytics due diligence tool (which is hosted on their webpage) to help them invest in early-stage companies;
- **GV**: everyone seems to know GV (formerly Google Ventures) is using AI and machine learning to inform their investment process, although almost no one knows exactly what they are doing and how;

- **Follow[the]seed**: a post-seed global algorithmic VC, they have developed two data-driven methodologies (one B2B and one B2C) to simply the investment process;
- **Right Side Capital Management**: with more than 800 pre-seed investments done so far, they make small investments ($100–$500 k at valuations of less than $3 M).

I have also heard Sunstone, e.ventures, and Nauta Capital use machine learning and analytics to some extent, but I could not find much information on them.

It is also worth to mention there are a few platforms (not VCs) that are looking at either helping investors or simply democratizing the investors' skills as much as possible. The first one is **Aingel.ai**, which has recently filed a patent for a machine learning system that scores startups and founders on the basis of a set of different variables. **PreSeries** is another fully automated solution to discover and evaluate startups, which it also has a voice interface (through Alexa).

15.4   Conclusions

The idea of a de-risked and unbiased selection of early-stage companies is certainly attractive, and I believe it deserves further studies and attention since it has the potential to increase the quality of companies that get funded (and founded as well).

Of course, having a better due diligence and decision process can favor investors but it does not solve all their problems. Whether the entrepreneur gets an investor's money does not depend, in fact, by the ability of the VC to do her due diligence but is rather driven by establishing a personal relationship and providing some additional value to the mere monetary contribution.

Moreover, it is hard to understand *ex-ante* what effects this class of models might have on the creation of new startups (maybe only a few clusters or group will be founded in the future to reflect the "success factors" identified by the AI models, in a sort of *adverse selection* phenomenon) and it is also quite cumbersome to disentangle and consider the effect of the VC additional value in computing the likelihood of success of a company.

Is being data-driven a value for investors? Certainly. Whether anyway a fully automatic VC can be created is still to be seen.

References

Belenzon, S., Chatterji, A. K., & Daley, B. (2017). Eponymous entrepreneurs. *American Economic Review, 107*(6), 1638–1655.

Bengtsson, O., & Hsu, D. (2010). *How do venture capital partners match with startup founders?* Unpublished working paper.

Cockburn, I., & MacGarvie, M. (2009). Patents, thickets and the financing of early-stage firms: Evidence from the software industry. *Journal of Economics and Management Strategy, 18,* 729–773.

Cumming, D. (2008). Contracts and exits in venture capital finance. *Review of Financial Studies, 21,* 1947–1982.

Da Rin, M., Hellmann, T., & Puri, M. (2013). A survey of venture capital research. *Handbook of the Economics of Finance, 2,* 573–648.

Eesley, C. E., Hsu, D. H., & Roberts, E. B. (2014). The contingent effects of top management teams on venture performance: Aligning founding team composition with innovation strategy and commercialization environment. *Strategic Management Journal, 35*(12), 1798–1817.

Ewens, M., & Marx, M. (2017). Founder replacement and startup performance. *The Review of Financial Studies, 31*(4), 1532–1565.

Ewens, M., & Rhodes-Kropf, M. (2015). Is a VC partnership greater than the sum of its partners? *The Journal of Finance, 70*(3), 1081–1113.

Franke, N., Gruber, M., Harhoff, D., & Henkel, J. (2006). What you are is what you like-similarity biases in venture capitalists' evaluations of start-up teams. *Journal of Business Venturing, 21* (6), 802–826.

Gloor, P. A., Dorsaz, P., & Fuehres, H. (2011). Analyzing success of startup entrepreneurs by measuring their social network distance to a business networking hub. In *Proceedings 3rd international conference on collaborative innovation networks coins* (pp. 8–10).

Gloor, P., Dorsaz, P., Fuehres, H., & Vogel, M. (2013). Choosing the right friends–predicting success of startup entrepreneurs and innovators through their online social network structure. *International Journal of Organisational Design and Engineering, 3*(1), 67–85.

Gompers, P., Kovner, A., & Lerner, J. (2009). Specialization and success: Evidence from venture capital. *Journal of Economics and Management Strategy, 18,* 817–844.

Gompers, P., Lerner, J., Scharfstein, D., & Kovner, A. (2010). Performance persistence in entrepreneurship. *Journal of Financial Economics, 96,* 18–32.

Groenewegen, G., & de Langen, F. (2012). Critical success factors of the survival of start-ups with a radical innovation. *Journal of Applied Economics and Business Research, 2*(3), 155–171.

Guzman, J., & Stern, S. (2014). *Nowcasting and placecasting entrepreneurial quality and performance*. MIT Working Paper.

Halabi, C. E., & Lussier, R. N. (2014). A model for predicting small firm performance: Increasing the probability of entrepreneurial success in Chile. *Journal of Small Business and Enterprise Development, 21*(1), 4–25.

Hoenig, D., & Henkel, J. (2015). Quality signals? The role of patents, alliances, and team experience in venture capital financing. *Research Policy, 44*(5), 1049–1064.

Hsu, D. (2004). What do entrepreneurs pay for venture capital affiliation? *Journal of Finance, 59,* 1805–1844.

Hsu, D. (2006). Venture capitalists and cooperative start-up commercialization strategy. *Management Science, 52,* 204–219.

Hsu, D. (2007). Experienced entrepreneurial founders, organizational capital, and venture capital funding. *Research Policy, 36,* 722–741.

Hsu, D., & Ziedonis, R. (2011). *Strategic factor markets and the financing of technology startups: When do patents matter more as signaling devices?* Unpublished working paper.

Inderst, R., & Mueller, H. M. (2009). Early-stage financing and firm growth in new industries. *Journal of Financial Economics, 93*(2), 276–291.

Lasch, F., Le Roy, F., & Yami, S. (2007). Emergence of high-tech ventures in France: How do regional, individual and organizational factors influence birth and sustainability of new firms? In J. Ulijn, D. Drillon, F. Lasch (Eds.), *Entrepreneurship, cooperation and the firm: The emergence and survival of high technology ventures in Europe* (pp. 187–218). Northampton, MA: Edward Elgar Publishing, Inc.

Lindsey, L. (2008). Blurring firm boundaries: The role of venture capital in strategic alliance. *Journal of Finance, 63,* 1137–1168.

Littunen, H., & Niitykangas, H. (2010). The rapid growth of young firms during various stages of entrepreneurship. *Journal of Small Business and Enterprise Development, 17*(1), 8–31.

Lussier, R. N., & Halabi, C. E. (2010). A three-country comparison of the business success versus failure prediction model. *Journal of Small Business Management, 48*(3), 360–377.

Mann, R., & Sager, T. (2007). Patents, venture capital, and software start-ups. *Research Policy, 36,* 193–208.

Marom, S., & Lussier, R. N. (2014). A business success versus failure prediction model for small businesses in Israel. *Business and Economic Research, 4*(2), 63–81.

McKenzie, D., & Paffhausen, A. L. (2017). *Small firm death in developing countries*. World Bank Policy Research Working Paper 8236.

McKenzie, D., & Sansone, D. (2017). *Man vs. machine in predicting successful entrepreneurs: Evidence from a business plan competition in Nigeria*. World Bank Policy Research Working Paper 8271: 1–66.

Miettinen, M. R., & Littunen, H. (2013). Factors contributing to the success of startup firms using two-point or multiple-point scale models. *Entrepreneurship Research Journal, 3*(4), 449–481.

Miloud, T., Aspelund, A., & Cabrol, M. (2012). Startup valuation by venture capitalists: An empirical study. *Venture Capital, 14*(2–3), 151–174.

Mueller, B., & Murmann, M. (2016). *The workforce composition of young firms and product innovation—Complementarities in the skills of founders and their early employees*. Centre for European Economic Research Discussion Paper 16–074.

Nahata, R. (2008). Venture capital reputation and investment performance. *Journal of Financial Economics, 90,* 127–151.

Nann, S., Krauss, J. S., Schober, M., Gloor, P. A., Fischbach, K., & Führes, H. (2010). *The power of alumni networks—Success of startup companies correlates with online social network structure of its founders*. MIT Sloan Research Paper: 4766–10.

Ng, W., & Stuart, T. (2016). *Of hobos and highfliers: Disentangling the classes and careers of technology based entrepreneurs*. Unpublished Working Paper.

Puri, M., & Zarutskie, R. (2012). On the lifecycle dynamics of venture-capital- and non-venture capital-financed firms. *Journal of Finance, 67*(6), 2247–2293.

Shane, S., & Stuart, T. (2002). Organizational endowments and the performance of university start-ups. *Management Science, 48,* 154–170.

Sheehan, P., & Sheehan, A. (2017). *The paradox of experience*. ETF Partners White Paper.

Sorensen, M. (2007). How smart is smart money? A two-sided matching model of venture capital. *The Journal of Finance, 62*(6), 2725–2762.

Sunesson, T. (2009). *Alma mater matters: The value of school ties in the venture capital industry*. Unpublished working paper.

Teng, H., Singh Bhatia, G., & Anwar, S. (2011). A success versus failure prediction model for small businesses in Singapore. *American Journal of Business, 26*(1), 50–64.

Tian, X. (2011). The causes and consequences of venture capital stage financing. *Journal of Financial Economics, 101,* 132–159.

Zacharakis, A. L., & Meyer, G. D. (2000). The potential of actuarial decision models: Can they improve the venture capital investment decision? *Journal of Business Venturing, 15*(4), 323–346.

Zacharakis, A. L., & Shepherd, D. A. (2001). The nature of information and overconfidence on venture capitalists' decision making. *Journal of Business Venturing, 16*(4), 311–332.

Zarutskie, R. (2010). The role of top management team human capital in venture capital markets: Evidence from first-time funds. *Journal of Business Venturing, 25,* 155–172.

Chapter 16
A Guide to AI Accelerators and Incubators

16.1 Definitions

The current startups landscape is incredibly messy. Venture capitalist, angels, incubators, accelerators, private equity funds, corporate venture capital, private companies, research grants. There are plenty of ways to get funded to start your own company—but how many of them are not simply '*dumb money*'? How many of them give you some additional value and really help you scale your business?

This problem is particularly relevant for emerging exponential technologies such as artificial intelligence, machine learning and robotics. For those specific fields, highly specialized investors/advisors are essential for the success of the venture.

This is why this chapter focuses on accelerators and incubators programs.

Since the edges are blurring, it is hard to find a commonly shared definition for accelerators and incubators. Hence, I will provide two different definitions, one a bit more from a practitioner's point of view, the other slightly more academic.

In the industry, the distinction between an accelerator and an incubator is simply related to the rationale for a company to join such a program. In other words, an incubator helps the entrepreneur in the **development of her idea**, while the accelerator focuses more on **growing the business**. The two programs have therefore two different goals and should be joined at a different stage of the startup lifecycle (Isabelle 2013).

If we look instead at a more rigorous detailed academic definition, it would be worth to have a look at Cohen (2013) and Cohen and Hochberg (2014). They actually define a startup accelerator as "a fixed-term, cohort-based program, including mentorship and educational components, that culminates in a public pitch event or demo day."

From this definition is clear that the authors looked at different traits to characterize and distinguish different programs from each other. The key features are indeed summarized in Cohen (2013).

Even though this academic definition clearly indicates thresholds and binary variables to identify different programs, it looks to me that—at least in the AI space—things are more complicated and actually it is really hard to define who is who (for help, check the brilliant review by Hausberg and Korreck 2017). Furthermore, the important question we should ask is not whether to call a program accelerator or incubator, but rather what is the real value brought to the entrepreneur.

16.2   Are They Worth Their Value?

If you are an entrepreneur, having so many different choices might make you wonder whether it might make sense to join one of those programs or not. And if you are an investor, a company, or anyone else looking at the space, you might start wondering if those programs suffer from an adverse selection problem: *good companies* go ahead with their feet while '*lemon' companies* that cannot get funded or get the ball rolling go into these programs.

16.2.1   *Entrepreneur Perspective: To Join or not to Join*

Unless you are already an experienced entrepreneur, the short answer is **yes, accelerators and incubators are worthy** (Hallen et al. 2016). Starting and running a company is something no university can teach you (no matter how many innovation workshops you take or entrepreneurial courses you attend) but it is grounded on real life experience. In this respect, accelerator programs are sort of full-time educational bootcamp in which you rapidly learn what you need to at least survive the first year. Whether then you are going to make it or not depends on how you transform that knowledge into the right actions.

Joining an accelerator is actually as reading a summary instead of the full book to do an exam: in this case, the full book would take you years to be read, while the summary takes a few months and can help you passing the exam. However, final graduation is a completely different thing.

Academic research, even if not unanimously (check this beautiful work by Yu 2016), seems to confirm with data the value of those programs (Hochberg 2015). Studies prove that accelerated companies reach milestones faster (Hallen et al. 2014), have a higher probability to raise further funding with respect to angel-supported startups (Winston-Smith and Hannigan 2015), and that have even spillover effects on the entire entrepreneurial ecosystem (Fehder and Hochberg 2014).

A warning though: even if some of those findings are true from a statistical point of view, there is a huge difference between different accelerators, and the **quality** of the program drastically impacts the positive effects for the startup.

16.2.2   *Investor Perspective: Should I Stay or Should I Go*

A good investor is basically the one who is able to:

(i) pick straight the winner and helping him become bigger and stronger;
(ii) pick a potential winner with the right things in place and helping him become successful.

The first case requires a lot of ex-ante work (due diligence) but not much after you invest. You simply seat down, relax and wait (it is not that simple actually, but let me go with this narrative for a second). The problem here is that there are few companies with these traits and everyone wants to invest in them, which considerably reduces the risk-return tradeoff.

The second case is instead more interesting and shows the real skills and contribution of the investor. It is also what it happens, most of the time and with exceptions, with companies coming out from accelerators and incubators program. These are companies that, for whatever reasons (lack of previous experience, no access to funding, etc.), might not have made it by themselves but are now in the game. Think of big success stories as Dropbox, for example.

So the question is: as an investor, should I invest in companies coming out from accelerator programs? Or am I buying a lemon?

The answer is 'simple', once again: yes, but mainly in those ones coming from excellent successful programs.

The proliferation of accelerators and incubators program made really difficult for investors to find real value in accelerated companies, especially for AI-related technologies and businesses. Good companies join accelerators for learning, mentoring and to get more exposure, all things as an entrepreneur you want to get from the best ones out there. And if good companies join an accelerator, the accelerator becomes more successful and attract better and better companies and founders on the next batches. It is a virtuous circle, which is creating a clear polarization in the industry, a positive skew distribution where very few programs deliver excellent results while the majority of them do not add any value (and in some cases are even detrimental) to the participants. In other words, I think there is a strong adverse selection problem in the accelerators/incubators space.

Of course, this is not a law of nature and does not imply that every company coming out from Techstars is going to become a unicorn (or the other way round). It is simply a rule of thumb to allocate a bit more efficiently your capital. If you are then able to spot out a potential winner in a low-level accelerator, chapeau, give yourself a pat on the shoulder because you did a very nice job.

16.2.3 *Accelerators Assessment Metrics: Is the Program Any Good?*

The common denominator of the two perspectives is that everything comes back to how good an acceleration program is. I have no particular experience in setting up or participating in an accelerator, so I do not know for sure the problems or the metrics on how to assess it. This is my interpretation (quite general with some sprinkle of AI somewhere), but feel free to comment below and tell me more about different metrics and aspects I should also consider:

(i) **Alumni network**: who are the alumni of the program? This base represents the 'customer base' of the accelerator, so check it out if includes big names. Do not be trapped by average valuations of the portfolio of the program: having one Dropbox and dozen of 'John Doe startups' does not make it a good accelerator, it simply makes it a lucky one (look at different stats, if you want to, e.g., median, variance, etc.);

(ii) **Raising the next round**: even though raising funds is not always a proof of business success, it is very often a good proxy for it. The more companies raise a further fund after the program, the better the program is;

(iii) **Raising a good next round**: same considerations as above, with the additional aspect that companies need to raise a specific amount of money. The more companies can reach their funding goal, the better the program is.
Be careful: evaluating an accelerator on the basis of the average amount of dollars raised is a huge mistake and only increments the already existing hype on AI;

(iv) **Survival rate**: the accelerators are set to provide entrepreneurs with tools and network to survive for at least 12 months (this is my view). The higher number of companies are still operating after one year, the better the accelerator was;

(v) **Exit**: ceteris paribus, if companies coming out from programs are obtaining higher valuation than their competitors, shortening the time-to-exit, or simply increasing the probability of an exit, it means that the accelerator did the job it was supposed to.
However, this point is controversial for at least two reasons: first, it is statistically hard to understand how an accelerator affects a final exit. Life is much more complicated than linking straight accelerator to a higher exit, but if all the companies coming out from a specific program obtain higher valuations with respect to their peers, we know for sure that there is some endogeneity there, even if we might not be able to identify the specific factors that make a business more successful.

Second, it depends on your view about business and what it means starting a company. Real visionary entrepreneurs do not start a company to sell it—they start something as it should run forever. An exit is somehow a defeat for some of them (there are exceptions, e.g., DeepMind), but the reality is that this class of entrepreneurs is disappearing. People start business nowadays with the idea in mind to sell out in 5 years to a specific buyer, or to use the technology developed to increase the salary base from $150 k (a normal salary in big tech companies in the US for an AI researcher) to $7 M (average amount got from acqui-hire in AI and machine learning sector).
I am not saying this is wrong and this is certainly what an investor wants, but it can invalidate the 'Exit' metric as one variable to track for accelerators' performance;

(vi) **Wider network**: a good accelerator has top-level mentors and knows how to engage them to be effective. It also has people behind who can really understand AI technologies and can help entrepreneurs with latest developments in research, or partners that can provide datasets for feeding neural nets.

16.3 A Comparison Between Accelerators

The following list is not exhaustive but it presents an overview of the landscape of the accelerators and incubators programs working specifically with or for AI companies. I am not going to discuss every single one in details (you can read some extra information on my material online) but simply providing a table that summarizes the key points of each program.

The Table 16.1 provides therefore a summary of all the information for 34 accelerators (only for those ones I could find information about). If you are interested in knowing why some accelerators don't disclose information, check the theoretical work of Kim and Wagman (2014). Please consider the value of the funding as expressed in accelerator's local currency (except for Creative Destruction Lab which is in US$) and the length of the programs expressed in months sometimes approximated if originally in weeks.

16.4 Final Thoughts

I tried to list all the accelerators I could find working specifically on AI, and I hope it will help someone out there. It looks clear to me now that:

(i) the on-going confusion between accelerators and incubators facilitated the creation of *mixed structures* which have characteristics of both the programs;

Table 16.1 AI accelerators and incubators

	Location	Relocation	Length (months)	Funding (k)	Mentorship	Equity	Available Spots	Goal	Tech 'n' Perks
AI Nexus Lab	NY	Yes	4	100	2 Technical full-time + 1 student fellow	8%	5	Develop the initial product and prepare for seed and series A	Access to NYU supercomputer; NVIDIA DGX-1; AWS; Google App Engine; Salesforce
Alexa Accelerator (Amazon)	Seattle	Yes	3	100+20	Techstars Network (Brad Feld, David Cohen, and others)	6%	N/A	Advancing voice-powered technologies	Office space; 400 perks worth over $1M
Allen Institute for Artificial Intelligence	Seattle	Yes	6	250	Access to 70 AI researchers and PhD holders on staff	N/A	2-5	Matching entrepreneurs and technologists to build AI startups together	Office space
Bosch DNA	Berlin / Bangalore	Yes (partial)	4	0	Bosch experts	0%	10	Establishing collaborations between Bosch India and startups	Incubation space; IP, legal, and PR consulting; Bosch platforms (SensorTech, XDK Kit, Bezirk)
Botcamp	NY	Yes	2.5	200	Twitter, Microsoft, Slack, Facebook, Amazon Alexa, Line, Kik, Discord, Skype, WeChat, and The New York Times	8%	N/A	Developing new conversational interface technologies	Howdy; Dexter; AWS;
Comet Labs	SF	Yes (optional)	3	0-75	1 member of investment team; Doug Bemis; Hasan Korre; and others	Variable	7	Transportation Lab is focused on AI mobility solutions. Open also to entrepreneurs with idea and no prototype	Data sets; space to pilot; B2B sales coaching; office space
Creative Destruction Lab (AI/ML track)	Toronto / Vancouver	No	9	0	Shivon Zilis, Barney Pell, and others	0%	50	Help founders with an MVP but not yet seed round	Osler (legal), KPMG (accounting),
Creative Destruction Lab (Quantum ML track)	Toronto	Yes (partial)	9	80	William Tunstall-Pedoe, Barney Pell, Geordie Rose, Sally Daub, Anthony Lacavera, Ted Livingston, James Cham, Matt Ocko, Lyon Wong, Steve Jurvetson	8%	40	Develop and support the world's largest batch of quantum machine learning startups	Cloud access to the D-Wave Sampling Service;
CyberLaunch	Atlanta	Yes	3	20-100	Sanjit Singh Dang, Galen Buckwalter, Chad Byers and others	7%	N/A	Acceleration with focus on building a scalable and repeatable business model	AWS; legal services; co-working space; 1 month of intensive technical training;
Data Elite Ventures	SF	Yes	3	150	Facebook, Netflix, Zynga	6%	5-10	Accelerate people and companies with already proven expertise in data science and big data (at least 5 years)	Office space
Deep Science Ventures	London	Yes	3+3	30	N/A	15%	Up to 100 scientists	A lab for UK's top talents to experiment and build the next generation of high-growth science based ventures from scratch (no idea or team required)	Office space
Founders Factory	London	Yes	6	30	CSC; Thomas Stone; etc.	7%	2-3	Co-creation/development accelerator for companies with already product and some traction	They offer their full time operations team to startups
H2 Ventures	Sydney	Yes	5+1.5	50	Stone & Chalk network plus 50 mentors from finance and tech companies	10%	8	Specialized in fintech and AI. Creating a working MVP and proof of market in 3-6 months and raising Series A in following 6-12 months.	They offer an extra 1.5 month of free desks before leaving (in addition to office space during the program).
IBM Alphazone	Israel	Yes	5	0	IBM, BD, and others	0%	5-8	Helping startups to build leading solutions for the enterprise market. The program focuses on post Seed & Round A funded companies with aim to create long term technology and business partnership with IBM worldwide	Office space; IBM SoftLayer / Bluemix Cloud Services; Free access to IBM technologies and tools (development licences);
Innovate8 Connect	Singapore	N/A	3	75	Singtel Group and Optus	0%	14	Programme that brings start-ups and Singtel together to create innovative solutions for business challenges faced by the Singtel Group	N/A
Kapsch Factory1	Vienna	No	6	150	20 mentoring sessions (remote & on-site)	0%	5	Acceleration program with a focus on future intelligent mobility solutions	Global industry events and fairs; PR and marketing support; financial support and strategic funding of global proof-of-concept projects
Microsoft Accelerator	Bangalore	Yes	4	0	150+ mentors (including Harvard Angels, RedBus, etc.)	0%	14	They target later stage startup to create MVC (minimum viable companies)	Microsoft Azure; Alumni program
Next AI	Toronto	Yes	5	50-200	IBM, Google, NVIDIA, Ajay Agrawal, Graham Taylor, Raquel Urtasun, and others	0%	N/A	Identifying talented teams with ambitious ideas and leverage Canada's leadership position in AI to provide them with the capital, mentorship, education and network to disrupt industries.	Office space; Legal and IP consulting; access to technology platforms (IBM, NVIDIA, Google) and datasets
NVIDIA Inception	Virtual	No	N/A	0	NVIDA (and its ecosystem)	0%	N/A	Helping startups during critical stages of product development, prototyping, and deployment (especially useful for who does intensive use of GPUs)	GPUs; credits; exclusive courses from DLI; marketing consulting
Play Labs	Cambridge (MA)	Yes	3	20	MIT Faculty and alumni mainly	6%	5-10	Only for MIT students and alumni to work on playful technologies	Office space
Rockstart	Netherlands	Yes	6	20	Jheronimus Academy of Data Science	6%	10	First Continental Europe accelerator for general AI startups	Office space
TechCode AI+	Mountain View (CA)	Yes	3	50	N/A	N/A	50	Programme for startups developing products incorporating artificial intelligence technologies and looking for access in China	Office space
The Hive	Palo Alto (CA)	Yes	Variable	1,500-2,000	Paul Maritz, Eduardo Castro-Wright, and many others	Variable	Variable	Co-creation studio to build and launch startups	Synapse; coworking space
Voicecamp	NY	Yes	2.5	200	Amazon, Microsoft, Slack, SoundHound, IBM, Greylock, Benchmark, KPCB, Bessemer, Homebrew, etc.	8%	8-10	Accelerator for early stage companies building voice-based products	Office space; credits for AWS, Digital Ocean, Twilio, Azure, SVB;
Winton Labs	London	Yes	3	5-15	Winton experts, Academic partners (London)	0%	5-6	For startups that are yet to receive a significant round of institutional funding in the areas of machine intelligence, forecasting, and innovative data	Office space;
YC AI	Mountain View (CA)	Yes	3	120	This list is infinite...	7%	Not fixed, but on average 90-100	Agnostic to the industry and would eventually like to fund an AI company in every vertical. Specific focus on Robot Factories	Office hours with ML engineers; Cloud credits for GPUs; Proprietary datasets and computing infrastructure (upon experiment success)
Zeroth AI	Hong Kong	Partial	3	20	Jaan Tallin (co-founder of Skype); Thomas Stone; Nathan Benaich; Azeem Achar; Jon Bradford and many others	6%	10	The programme's success is defined by each startup's ability to find and grow users for its product.	Office hours; office space;

(ii) quality matters (not all the accelerator are equals). You get different value from different ecosystems even if the offer is the same on paper. Joining an accelerator in this list is also not a guarantee of success, and of course, there are many other excellent programs worldwide that can maybe work much better than some of the ones I showed above.

The motif, though (and my personal believe at this stage of AI development), is that *specialized investors and accelerators* can do a much better job in understanding and helping companies leveraging these exponential technologies.

There is also something else emerging from the list: there are really few AI accelerators/incubators in Silicon Valley proportionally speaking, although the common expectation would be to find most of them in the American entrepreneurial district.

My guess is that, in reality, from a pure cost-benefit perspective, the Bay Area is not the best place to start a company.

It is the best place though to expose the startup to a larger market, investors and public acknowledgement.

This does not imply that being in Silicon Valley makes no sense, but rather the opposite. I actually see shaping an emerging pattern in Silicon Valley, the same one that characterized in the past 30 years the pharmaceutical and movie industries. The pharma industry, for example, moved from being a large industry where the same company did the research (expensive), developed the molecules (expensive) and eventually commercialized the final product (cheap and with good margins), into a two-ways sector where biotech companies took the higher risk of developing experimental molecules while big pharma corporations oversaw FDA regulation approval and market launch.

Of course, it is a bit more complicated than that, but the main message is that the **sector self-specialized and assigned to each class of players what they knew how to do more efficiently** (research for biotech and commercialization for pharma companies).

In the same way, it will make sense probably to develop companies in other countries (where the real cost of starting up is much lower) to eventually land in California only once ready to either scale, raise larger rounds of financing or massively go to market.

A final interesting thing I noticed, which might be useful to some entrepreneurs: it is coming out the new concept of '*specialized co-working space*', and we have something focusing on AI called **RobotX Space** in multiple cities (Silicon Valley and Asia). I have never been there (but hopefully I will in the future) but I think that it makes a lot of sense to create technology hubs like this one. This model might, in the future, even undermine the business models of accelerators and incubators.

References

Cohen, S. (2013). What do accelerators do? Insights from incubators and angels. *Innovations, 8*(3/4), 19–25.

Cohen, S., & Hochberg, Y. V. (2014). *Accelerating startups: The seed accelerator phenomenon*. Working paper.

Fehder, D. C., & Hochberg, Y. V. (2014). *Accelerators and the regional supply of venture capital investment*. Working paper.

Hallen, B. L., Bingham, C., & Cohen, S. (2014). Do accelerators accelerate? A study of venture accelerators as a path to success. *Academy of Management Annual Meeting Proceedings*.

Hallen, B. L., Bingham, C., & Cohen, S. (2016). Do accelerators accelerate? *The Role of Indirect Learning in New Venture Development*. Available at SSRN: https://ssrn.com/abstract=2719810.

Hausberg, J. P., & Korreck, S. (2017). *A systematic review and research agenda on incubators and accelerators*. Available at SSRN: https://ssrn.com/abstract=2919340.

Hochberg, Y. V. (2015). Accelerating entrepreneurs and ecosystems: The seed accelerator model. In J. Lerner, & S. Stern (Eds.), *Innovation policy and the economy* (Vol. 16). National Bureau of Economic Research.

Isabelle, D. A. (2013). Key factors affecting a technology entrepreneur's choice of incubator or accelerator. *Technology Innovation Management Review*, 16–22.

Kim, J. H., & Wagman, L. (2014). Portfolio size and information disclosure: An analysis of startup accelerators. *Journal of Corporate Finance, 29,* 520–534.

Yu, S. (2016). *How do accelerators impact the performance of high-technology ventures?* Available at SSRN: https://ssrn.com/abstract=2503510.

Winston-Smith, S., & Hannigan, T. J. (2015). *Swinging for the fences: How do top accelerators impact the trajectories of new ventures?* Working paper.

Appendix A
Nomenclature for Managers

Relational database management system (RDBMS): structured data in predetermined schema (*tables*), scalable vertically through large SMP servers, or horizontally through clustering software. These databases are usually easy to create, access, and extend. The standard language for relational database interoperability is the **Structured Query Language (SQL)**.

Non-relational database: database that does not store data into tables, but made them accessible through special query APIs. The standard language used is **Not Only SQL (NoSQL)**: it does not present a fixed schema, it uses BASE system to scale vertically (basically available, soft-state, eventually consistent), and sharding (horizontal partitioning) to scale horizontally. Examples are **MongoDB** and **CouchDB** (they differ mainly because in MongoDB the main objects are documents, while in CouchDB are collections, which in turn contain documents). NoSQL commonly used **JavaScript Object Notation (JSON)** data format (**BSON** in MongoDB—binary JSON), and it mainly works through **Key Value Store (KSV)**, i.e., a collection of different unknown data types (while a RDBMS stores data into table knowing exactly the data type).

Hadoop: open source software for analyzing huge amount of data on a distributed system. His primary storage is called Hadoop distributed file system (**HDFS**), which duplicates the data and allocates them in different nodes. It has been written in Java. It is a core technology in the big data revolution and stores data into their native raw format, and it can be used for several purposes (Dull, 2014), such as a simply data staging or landing platform complementary to the existing EDW (as an enterprise data hub, i.e., EDH), or managing data (even small), transforming those into a specific format in the HDFS and sending them back to the EDW, lowering thus the costs while increasing the processing power. Furthermore, it can integrate external data-sources and archive data (both on-premises or into the cloud), and reduce the burden for a standard EDW.

MapReduce: software for parallel processing huge amount of data.

Flume: service to gather, aggregate, and move chunks of data from several sources to a centralized system.

Cassandra: an open source database system for analyzing large amount of data on a distributed system. It is characterized by a high performance and by a high availability with no single point of failure (i.e., a part of system that if fails stop the whole system). It fosters data denormalization, which means grouping data or adding redundant information, in order to optimize the database performance.

Distributed System: Multiple terminals communicating between them. The problem is divided in many tasks, and assigned to each terminal. It is a highly scalable system as further nodes are added.

Google File System: proprietary distributed file system for managing efficiently large datasets.

HBase: an open source non-relational database (column-oriented) developed on a HDFS. It is very useful for real time random read and write access to data, as well as to store sparse data (small specific chunk of data within a vast amount of them). The relational counterpart is called **Big Table**.

Enterprise Data Warehouse (EDW): system used for analysis and reporting that consists of central repositories of integrated data from a wide spectrum of different sources. The typical form of an EDW is the **extract-transform-load (ETL)**, that is the most representative case of *bulk data movement*, but other three important examples of these systems are **data marts** (i.e., a subset of the EDW extracted out in order to address a specific question), **Online analytical processing (OLAP)**—used for multidimensional low-frequency analytical query—and **Online transaction processing (OLTP)**—used rather for high volume fast transactional data processing. The wider system that includes instead a set of servers, storage, operating systems, database, business intelligence, data mining, etc. is called **data warehouse appliance (DWA)**.

Resilient Distributed Datasets (RDD): logical collection of data partitioned across machines. The most known examples is **Spark**, an open source clustering computing that has been designed to accelerate analytics on Hadoop thanks to the multi-stage in-memory primitives (that are basic data types defined in programming languages or built it with their support). It seems to run 100 times faster than Hadoop, but its disadvantage is that it does not provide its own distributed storage system.

Hive: additional example of EDW infrastructure that facilitates data summarization, ad-hoc queries, and specific analysis.

Pig: platform for processing huge amount of data through a native programming language called Pig Latin. It runs at the same time sequences of MapReduce.

Programming language: is a formal constructed language designed to communicate instructions to a machine. The main ones for data science applications are Java, C, C++, C#, R, and Matlab. Scala is another language that is becoming extremely popular right now.

Scripting Language: is a programming language that supports scripts, which are piece of codes written for a run-time environment that interpret (rather than

compile) and automate the execution of tasks. The main ones in big data field are Python, JavaScript, PHP, Perl, Rub, and Visual Basic Script.

Data Mart: is a subset of the data warehouse used for a specific purpose. Data marts are then department-specific or related to a single line of business (LoB). The next level of data marts is the **Virtual Data Marts**, i.e., a virtual layer that create various views of data slices—in other words, instead of physically creating a data mart, it just takes a snapshot of them. The final evolution is instead called **Data Lakes**, which are massive repositories of unstructured data with an incredible computational capability. Hence, data marts physically create repositories (slices) of data, virtual data marts leave the data where they are and create virtual constructs— reducing the cost of transferring and replicating them—while data lakes work as the virtual data marts but with any kind of data format.

Appendix B
Data Science Maturity Test

The following questionnaire provided could help managers to grasp a rough idea of the current data stage of maturity they are facing within their organizations. It has to be integrated with deep conversations and meetings with the big data analytics (BDA) staff, the IT team, and supported by solid researches.

(1) What is your investment level in BDA capabilities?

    1. Absent. We don't have money for big data
    2. A small budget is allocated when positive quarters in core activities allow us to do that
    3. A modest funding scheme is in place
    4. We invested a good percentage of our revenues in BDA in the last year, and we will keep investing because it is part of our company's vision.

(2) What the executives' support to analytics capabilities?

    1. Neither IT nor business think BDA is useful to the business
    2. Only IT managers support it because they are interested in the technological challenge
    3. Business managers see the hidden value in data and support BDA projects
    4. Both IT and business executives believe in BDA potential.

(3) What is your current stage of working with data?

    1. We will start using data in the future if needed
    2. We have a good idea of what business questions we could solve with data in my company
    3. We take action using analytics

    4. We are automating analytics the most we can, and we believe is a competitive factor that gives us benefits we are able to communicate frequently to top management and shareholders.

(4) Your analytics team is:

    1. Inexistent
    2. Acquired from outside at the moment
    3. We have some senior scientist that has been recruited, but we are now growing the team internally by training
    4. An independent sustainable group and function within the company.

(5) Your company's culture is:

    1. Intolerant—especially for failure concerning new analytics, methodologies, and technologies
    2. Variegated—it is half-half made by old-style professionals and geeks
    3. Collaborative—people are willing to work together and share.
    4. Creative—innovation is valued and we are encouraged and monetarily compensated for our original shared contributions.

(6) How your data science team is connected to the company hierarchy?

    1. We only have some analysts with small tasks, who deliver the outcomes to their direct managers on a weekly/monthly basis
    2. The data team is leaded by a business head, and their contribution is continuously marginally positive
    3. Our data scientists are tight to our data warehouse and data management teams, and they constantly interconnected with the business side
    4. They are autonomous and do not seat in the same building of the operations function. They are allocated in a Centre of Excellence.

(7) The internal data policy is:

    1. Fairly poor, we do not need it
    2. Metadata definitions and BDA policy are well-established
    3. We have a BDA policy that we constantly monitor and we have a security policy for any data forms
    4. We have a BDA and security policies, and we anonymized all the relevant data to protect our clients and partners' privacy.

(8) The data in your company are:

    1. Stored in silos
    2. We prioritized the data to be used within our organization, and they are internally shared

3. Many different data sources are integrated for our analysis, and we take care of data quality through a meticulous goodness assessment based either on the final use or the type of data we will exploit
4. We have integrated BDA technologies into our systems, we store our data on a cloud, and we often use them for mobile applications.

(9) When your company looks at your BDA capabilities:

1. It sees mainly a sunk-cost, i.e., the cost of storing, maintaining, protecting and analyzing these datasets
2. We know data have value and we understand both the data cost and data competitive advantage, but we are definitely overwhelmed
3. We are rationalizing our data storage and usage abilities, because we understood that not everything is either pertinent or meaningful
4. We have an efficient process for data aggregation, integration, normalization and analysis, and we can manage easily any amount of inflowing data.

(10) Your firm is currently using:

1. Relational Database and Internal data
2. Data marts, R or Python languages, and public data
3. NoSQL database, Hadoop and MapReduce, and we use external data, sometimes also unstructured
4. Highly unstructured data, APIs, and a Resilient Database.

Once each single question has been answered, it is simple to obtain a rough measure of the data maturity stage for a certain company. For each answer indeed, it has to be considered the number associated to that answer, and then it is enough to sum up all the numbers obtained in this way. So, for example, if in the third question the answer is "we take action using analytics", the number to be considered is 3, since it is the third answer of the list.

Finally, the score obtained should range between 1 and 40. The company will then belong to one of the four stages explained in Table 2.1 accordingly to the score achieved, that is explained in the Table B.1:

Table B.1 Data science maturity test classification

	Primitive	Bespoke	Factory	Scientific
Score	10–15	16–25	26–35	36–40

Appendix C
Data Scientist Extended Skills List
(Examples in Parentheses)

Programming (R, Python, Scala, JavaScript, Java, Ruby, C++).

Statistics and Econometrics (probability theory, ANOVA, MLE, regressions, time series, spatial statistics).

Bayesian Statistics (MCMC, Gibbs sampling, MH Algorithm, Hidden Markov Model).

Machine Learning (supervised and unsupervised learning, CART).

Mathematics (Matrix algebra, relational algebra, calculus).

Big Data Platforms (Hadoop, Map/Reduce, Hive, Pig, Spark).

Text mining (Natural Language Processing, SVM, LDA, LSA).

Visualization (graph analysis, social/Bayes/neural networks, Tableau, ggplot, D3, Gephi, Neo4j).

Business (business and product development, budgeting and funding, project management, marketing surveys).

Algorithms (SVM, PCA, GMM, K-means, Deep Learning).

Optimization (linear, integer, convex, global).

Simulations (Monte Carlo, agent-based modeling, NetLogo).

Structured Dataset (SQL, JSON, BigTable).

Unstructured Dataset (text, audio, video, BSON, noSQL, MongoDB, CouchDB).

Multi-structured Dataset (IoT, M2M).

Data Analysis (feature extraction, stratified sampling, data integration, normalization, web scraping).

Systems Architecture and Administration (DBA, SAN, cloud, Apache, RDBMS).

Scientific approach (experimental design, A/B testing, technical writing skills, RCT).

Appendix D
Data Scientist Personality Questionnaire

The terminology used to classify into 16 subcategories the different kind of data scientists is given by the two-entry matrix exhibited in the Table 7.1. The terminology can be sometime misleading if related to the Keirsey Temperament Sorter (KTS), and this is why it is necessary to specify that the only categorization borrowed from KTS framework is the broader one, i.e. the Artisan-Idealist-Rational-Guardian partition. Every sub-category has instead to be taken as newly generated.

Here it follows the personality test to sort data scientists into a specific box. It is composed by 10 questions, and for each one a single answer has to be provided. This test is not a professional temperament test to fully understand individuals' personality, but it is more a quick tool for managers to efficiently and consciously allocate the right people to the right team.

(1) When you start working on a new dataset

    a. You start exploring immediately and querying the data
    b. Plan in advance how to tackle it
    c. You spent time in understanding the data, where they come from, and their meaning
    d. You identify a research question quickly, and focus on designing the a new improved method for analyzing your data.

(2) In your team, people count on you for your

    a. Troubleshooting ability
    b. Organizational skills
    c. Capacity to reduce the problem complexity
    d. Strategic approach and conceptualization of the problem.

(3) When facing a new data challenge, your first thought is

    a. Is what I am doing impactful and relevant?
    b. When do I have to deliver some results?
    c. How this challenge can make me better?
    d. What I can learn from this dataset?

(4) In a data analysis, which is the most important thing to you

    a. Results, no matter how you do achieve them, what strategy or technology you do employ
    b. To achieve a result in the correct way and with the right process or technology
    c. Attaining significant results in an ethical manner
    d. Reaching the outcomes through an accurate, replicable, and efficient procedure.

(5) If you have finished your assigned today's work, you would

    a. Focus again on your analysis and try to find alternative and innovative way to achieve your final goal
    b. Start with something else, even if this may mean to stay longer at your desk
    c. Help a colleague in difficulty with his analysis
    d. Give suggestions and highlight weaknesses in your colleagues' works for the sake of the team and business development.

(6) If you would have some spare time during your daily work, you would prefer to

    e. Optimize existing technology for the whole company
    f. Improve your analysis
    g. Try to derive new insights from your previous analysis
    h. Understanding how to maximize the value of your analysis.

(7) It is your data-dream of

    e. Speaking about data with only engineers and IT team
    f. Teaching data related contents
    g. Engaging with people who do not know anything about data science
    h. Persuading and convincing the business team of the big data opportunity

(8) You prefer to work with

    e. Huge amount of structured data
    f. Any kind of data that challenge me
    g. Behavioral or social media data, or any unusual data
    h. No data in particular.

(9) If you would quit tomorrow your data science job, you would prefer to become

    e. An IT manager or software engineer
    f. A professor
    g. A consultant
    h. An entrepreneur.

(10) What characteristic of big data you value the most

    e. Volume
    f. Velocity
    g. Variety
    h. Value.

Once each question has been answered with a single reply, the result is given by pairing the reply chosen more often within the first five questions (a–d) with the answer that appears more often in the last five (e–f), as shown in the following table. So, if for instance in the first five questions *b* emerges as predominant answer, while in the last five *f* is the median, the person considered is a *Cruncher* (Table D.1).

Table D.1 Data scientist personality classification

Archetype/ personality	Artisan	Guardian	Idealist	Rational
Technical	*Gardener*: A–E	*Architect*: B–E	*Evangelist*: C–E	*Wrangler*: D–E
Researcher	*Alchemist*: A–F	*Cruncher*: B–F	*Champion*: C–F	*Groundbreaker*: D–F
Creative	*Trailblazer*: A–G	*Catalyst*: B–G	*Visionary*: C–G	*Warlock*: D–G
Strategist	*Babelian*: A–H	*Mastermind*: B–H	*Advocate*: C–H	*Fisherman*: D–H