

Introduction

“Artificial intelligence will reach human levels by around 2029. Follow that out further to, say, 2045, we will have multiplied the intelligence, the human biological machine intelligence of our civilization a billion-fold.”

—Ray Kurzweil

Evolution is the story of tension and mutation. A difference between aspirations and available resources creates tension while changing needs and advancing technology drive mutations. Successful mutations thrive until eventually they, too, are replaced by more successful variants. By its very definition and design, evolution is a slow process. However, every once in a while, evolution is known to change pace, resulting in rapid strides that suddenly leave the standard way of life redundant and usher in a new era.

We live in an increasingly digital world – one in which we are inundated by IoT devices, connectivity, data, applications, and new experiences every moment. Our economy and our way of life are rapidly transforming into a digital one for most parts. This change is requiring businesses to rapidly reinvent themselves in this new world order, not only to thrive but sometimes just to survive.

In very recent times, Artificial Intelligence (AI) has taken center stage in the digital space. AI has existed as a discipline for more than 60 years now; its recent rejuvenation is driven by the advances in digital capabilities around IoT, big data management, cloud computing, and communication technologies. As per a recent McKinsey study,⁽¹⁾ by the end of 2030, the impact of AI is expected to be about \$13 trillion with over 70% of companies impacted by AI. An Accenture study on the subject talks about a similar impact and interestingly identifies innovation diffusion as a key benefit in addition to productivity-related benefits. A Forbes study quoting a similar influence cites AI as the “new electricity” already driving a \$2 trillion impact on the economy. Whichever study we refer to, the effect of AI on the economy and our lives is undeniable. While these numbers are pre-COVID-19 crisis, they still hold. While the timing may shift as the world

starts returning to normalcy, there is no doubt about the extent of the impact. As we will see in several later chapters, COVID-19 has increased the need for AI in every aspect of our lives.

The promise of AI was rooted in the desire to simulate human cognition in the machines as they started becoming more pervasive in the first half of the 20th century. While the academic community and selected industry researchers were successful in developing the mathematical and statistical foundation to solve several AI problems, lack of data and computational power limited the practical implementation of AI on a large scale. These recent digital technological advances allow us to now connect to a very large population of people and things to collect huge volumes and varieties of data, manage and process that data inexpensively and effectively, and apply the AI-based algorithmic analytics to create new insights and drive new experiences and outcomes.

Let us first understand why AI even matters. In the last couple of years, we have already moved from an information age to an algorithmic age.⁽²⁾ A few factors have driven this transition:

- Human sensory functions and the human brain have significantly more things to process today. This means that our attention spans are reducing and we also need capabilities to react fast to new inputs and events. We become more effective if along with information we are given more actionable insights. Consequently, we prefer to engage with different types of businesses and services that give us more actionable insights along with the information.
- Similarly, businesses are also flooded with significantly more information than they were ever exposed to and ready to deal with. Information is now coming from the connectivity of their devices and processes, inputs from users, and other types of transactional data that is getting generated every second. On top of that, they are in a complicated competitive environment that is rapidly transforming as technology is reducing the barriers of historical knowledge, access, and infrastructure required by businesses to be successful. Businesses now have to compete for users' attention every second and use that opportunity to demonstrate better value than others. Businesses also need to quickly innovate around their business models and services to avoid being disrupted.

Traditional business practices and underlying technology infrastructure that were based on predefined rules do not allow one to respond to these types of rapid synthesis and reaction scenarios. It is not possible to know every possible scenario and plan for it, so we need more machine-driven cognitive capabilities that require AI to now become more pervasive, hence the transition from the information age to algorithmic age. In his Future of Life Institute article, AI researcher at the University of Louisville Roman Yampolskiy said, "AI makes over 85% of all stock trades, controls operation of power plants, nuclear reactors, electric grid, traffic light coordination and . . . military nuclear response. . . . Complexity and speed required to meaningfully control those sophisticated processes prevent meaningful human control. We are simply not quick enough to respond to ultrafast events such as those in algorithmic trading and . . . in military drones."

While AI brings a lot of new possibilities, it also brings new problems because businesses have to now reinvent themselves in this new world order. Traditional economic models and management practices are being disrupted. Companies that have considered AI as only a technology enabler for incremental improvements have missed the mark and there is increasing

realization among business leaders about that. Those who have taken a more holistic transformative approach are emerging as the leaders in their space in this new world order. The largest technology companies of the world like Microsoft and Google are pivoting with more AI centrality; financial services companies like JPMC are using AI to renovate their client engagement; a traditional industrial giant like Schindler is driving rapid growth using digital and AI; a new age environment management company like Bee'ah is basing its future on AI; and small progressive nations like UAE have made AI a national agenda with possibly being the only country with a separate minister of AI.

The recent COVID-19 crisis has amplified the need for us to focus on digitization and AI. Our world rapidly changed in the first three months of 2020. From being highly integrated and interdependent, we moved into an era of isolation, containment, fear, and remote working in no time. A previously vibrant global economy is seeing one of its worst crises ever; many companies will not make it through this crisis. Organizational resilience in being able to meet the challenges of an unpredictable future is key to survival in this environment. Analytics today is a core capability for achieving such resilience. AI has also been incredibly useful in working through the crisis, forecasting the rate and direction of infection spread, and helping with decision support for containment strategy effectiveness and even with research around the vaccine.

While AI did see many springs and winters over the past 70 years of its existence, it is now firmly intertwined with everything we do. Let us begin our journey into the exciting world of AI.

EVOLUTION OF AI

This book is about the future, but we will spend a little time discussing the past because understanding evolution gives us context and helps us appreciate the path to the future. In this section, we have divided the evolution into multiple eras; these are our choices, they are not industry standards.

Since the beginning of civilization, humans have been ruminating about recreating human-like capabilities in machines. Starting with Greek mythology, there are references to machine-men. Fiction literature is littered with examples of artificial intelligence for a couple of hundred years now. Mary Shelley's *Frankenstein* (1818), Samuel Butler's *Darwin Among the Machines* (1863), and Karel Čapek's *R.U.R. Rossum's Universal Robots* (1920) are some examples of publications where concepts of intelligent devices and robots were discussed. Beyond literature, these ideas also showed up in motion pictures and television from the early days. Arguably, the 1927 German science-fiction film *Metropolis* is the first movie to depict a robot. A very popular show from the 1960s in the United States, *The Jetsons*, also very accurately depicted a future world. While the show was set in 2062, a number of the cool technologies shown in that are already part of our daily lives; these include video calls, robotic vacuum cleaners, robotic assistants, tablet computers, smartwatches, drones, holograms, flying cars, flat televisions, jetpacks, and many more. The trend of depicting a technology-enriched future has continued to the present day. We often take inspiration in life from literature and fiction. That can be said to be true for AI as well.

The 1950s – The Nativity Era

The birth of AI as a formal discipline of study and practice happened in the 1950s. Alan Turing, John McCarthy, Marvin Minsky, Allen Newell, Nathaniel Rochester, Claude Shannon, and Herbert A. Simon are considered as the founding fathers of AI. Norbert Wiener laid the foundation for cybernetics, Claude Shannon conceptualized information theory, and Alan Turing described thinking machines. The confluence of these ideas led to the development of AI. The term *artificial intelligence* was first coined by John McCarthy. Dartmouth Summer Research Project on Artificial Intelligence in the summer of 1956 was the event where AI was first formalized. It started as a proposal made several months earlier by McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon where the term *AI* was first used. The proposal states,

We propose that a 2-month, 10-man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed based on the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.⁽³⁾

But AI did not originate suddenly because of the academic research from the 1940s and 1950s. The roots of AI lie in formal reasoning, which has a very long history of over 2,000 years. References to formal reasoning can be found in texts from ancient Greece, India, and the Chinese. It is interesting to note that many early mathematicians were also philosophers. German mathematician Gottfried Leibniz was one of the first to suggest that human reason could be explained as mechanical calculations. *Principia Mathematica* by Alfred North Whitehead and Bertrand Russell published in 1910 is one of the seminal works that laid the foundation for modern mathematics and logic. For a very long time, we humans have been making progress in our quest to create intelligence artificially.

Seventy years back, in the decade of the 1950s, hype cycles did not exist and technology took a long time to transition from the academic realm to practical industry implementations. AI faced the same fate. In its decade of birth, the world saw foundational concepts around AI getting established. But there was a lot of optimism generated in this decade because people could now realistically see computers (machines) undertaking tasks, albeit simple ones, that so far only existed in the fiction realm. One important development in that decade was the invention of one of the earliest programming languages, Lisp, by John McCarthy in 1958. Its significance was realized in later decades as it became the mainstay for AI programming for a long time.

The 1960s – The Foundation Era

The optimism and energy of the 1950s led to the development of a lot of fundamental theories in the 1960s that form the basis of AI even today. Many advancements gave AI the legs to move from a theoretical possibility to practical applicability. Here are some of the highlights:

- **Generic systemic and programming approach to solving any problem.** Development of General Problem Solver technique. At the closing of the previous decade, Herbert A. Simon, J. C. Shaw, and Allen Newell wrote the first program for unconstrained problem-solving.
- **Using math to solve AI problems.** The introduction of Bayesian methods for inductive inference and prediction. Ray Solomonoff took the original theorem of Thomas Bayes and improvements to it made by Pierre-Simon Laplace and developed it further to meet the needs of AI.
- **Introduction of robotics.** Unimate was the first industrial robot that worked on a General Motors assembly line to transport die castings from an assembly line and welding these parts on autobodies. This was a complex task, and with the machine doing it, it increased the safety of workers who were routinely injured in executing that task. Toward the end of the decade, Shakey the Robot emerged as the first general-purpose robot that could analyze commands and decompose them into smaller executable components before carrying out the commanded tasks.
- **Mathematically modeling past data to predict the future.** Introduction of Machine Learning (ML). The roots of ML go back more than 200 years when Thomas Bayes created the Bayes Theorem. In the preceding decades, a lot of discovery around ML happened. But this is the decade when it formally started taking structure post the invention of Perceptron, the introduction of Bayesian methods in AI, and the development of other algorithms.
- **Computers were able to understand and process human language.** Introduction of Natural Language Processing (NLP). Joseph Weizenbaum, professor of the MIT Artificial Intelligence Lab, built ELIZA, the first NLP program that could converse about any topic in English.

Buoyed by these developments, the decade ended with further optimism. MIT AI scientist Marvin Minsky predicted that “in from three to eight years we will have a machine with the general intelligence of an average human being.” But that was not to be and AI will hit its first major brakes in the next decade.

The 1970s – The First Winter Era

From the mid-1970s to early 1980s is known as the First Winter of AI when the progress was halted by technological challenges, increasing skepticism, lack of funding, and very limited progress in real-world applications.

That era was also the early days of development of computers. There was not enough computing power available for many of the algorithms to work in any meaningful way. For example, the promise of NLP was dampened because the vocabulary could not be expanded beyond 20 words or so. Pattern recognition and image processing could not progress further for similar reasons. The capabilities of robots remained stagnant. While in the previous decade many academics went on celebrating the promise of quick progress in AI, in this decade a different group of academics were very vocal critics of the possibilities of AI and its inability to solve problems at a meaningful scale. Most of the funding for AI research so far was sponsored by government bodies, a number of them associated with the military. Consequent to the lack of progress and ensuing criticism, most of the funding got pulled.

However, this was not a completely lost decade. There was a lot of theoretical research and publishing that happened in this decade around visual perception, natural language processing, and various algorithmic approaches. Three very important things did happen in this decade that had a significant impact in the years to come:

- Herbert A. Simon won the Nobel Prize in Economics for his theory of bounded rationality, one of the cornerstones of AI.
- The concept of schemas and semantic interpretation of data was introduced by Marvin Minsky in 1975.
- The world's first autonomous vehicle, albeit a basic one but fully computer-controlled, was built by Hans Moravec in the Stanford AI Lab.

The pessimism and slowdown in the last few years of the decade helped AI jumpstart in the following decade.

The 1980s – The Resurrection Era

AI got a new lease of life in the 1980s and this decade became crucial for the advancement of many capabilities that will propel AI forward. Progress was seen both on the hardware and algorithmic side. Furthermore, government support for AI research started to flow through again. Here are some of the highlights from this era:

- **Rise of expert systems to solve specific problems.** By mimicking experts in specific domains and using algorithms to solve more well-defined problems, AI again started proving its value.
- **Machine learning got revived.** The introduction of the Backpropagation technique enabled error correction in prediction models.
- **Initiation of the Fifth-Generation Computer Systems project.** The Japanese government effort started the introduction of massively parallel computing to solve AI problems and take the machines closer to human reasoning levels.
- **Computing power-enabled ANN algorithms.** New electronic transistors and very large-scale system integration (VLSI) development for integrated chips allowed the processing power of computers to dramatically rise, enabling processing-intensive algorithms like artificial neural networks.
- **Development of knowledge systems.** Starting with Cyc, many advanced knowledge management systems using AI started getting developed; the foundation was laid for future projects like Deep Blue that catapulted AI to the next level.

This decade accomplished many things to give AI the ability to scale in solving complex real-world issues and break out of the academic realm.

The Brief Interlude

The years 1987 to 1993 saw another slowdown of the AI-wave largely driven by the inability of the hardware and computing capabilities to keep up with the algorithmic progress. Consequently,

the broader impact was not felt and funding was again cut. However, this did not last for too long. The research community retreated in the background and operated behind the shadows to come back with some impressive accomplishments in the 1990s.

The 1990s and 2000s – The Second Revival

Finally, during this era, AI broke out of the academic world and started becoming more mainstream. This was greatly aided by the improvements in computing power and the development of distributed systems. A lot of the algorithmic development, originally done for AI, finally proliferated the business world for data mining, search (like Google), robotics, medical diagnostics, financial services, and other industrial applications. In the 1990s, there were many other notable games in which AI-powered computer programs won. AI program DART was deployed for military purposes in the First Gulf War. In the first decade of the next century, a lot was achieved in robotics and autonomous driving. Many disciplines of AI like machine learning, intelligent tutoring, case-based reasoning, multi-agent planning, scheduling, uncertain reasoning, data mining, natural language understanding and translation, vision, virtual reality, games, and other topics got an opportunity for real-life demonstrations.

This was the time when the world was first consumed by the Y2K crisis, then by the Internet, and thereafter the massive boom in e-commerce and dot-com companies. So, while AI did not get much publicity, it kept making progress in the background. This helped because there were no inflated expectations to meet, yet researchers were able to contribute to solving business problems. Two very notable things did happen in this era that eventually led to the current hype around AI:

- Deep Blue defeated chess champion Gary Kasparov in 1997 and created a huge buzz around the power of computers and algorithms.
- Multi-agent systems started to mature with new concepts from decision theory to create intelligent agents; we shall discuss multi-agent systems a few times later in the book.

The world celebrated 50 years of AI in 2006. AI gained a lot of ground during its second revival. It took a bit of time, but after its first 50 years of existence, AI started to gain significant traction.

The 2010s – The Renaissance Era

The decade of 2010s will be known as the time when AI took center stage in technology-led transformations. The primacy of AI can be attributed to the gains made by big data, cloud, and IoT in the preceding years as well as the continued development of AI in the background. Advances in big data, cloud, and IoT solved two major issues that had plagued the progress of AI: enabled enough data and ensured enough compute capabilities. Deep learning became feasible and amplified the value of ML and AI more broadly, expanding its utility to many more problem spaces than before. Major technology players like Microsoft, Google, Amazon, and others made AI their main agenda in this period. Since the developments are recent, we will not get into more detail here.

Figure 1.1 depicts the evolution of AI over the past seven decades.

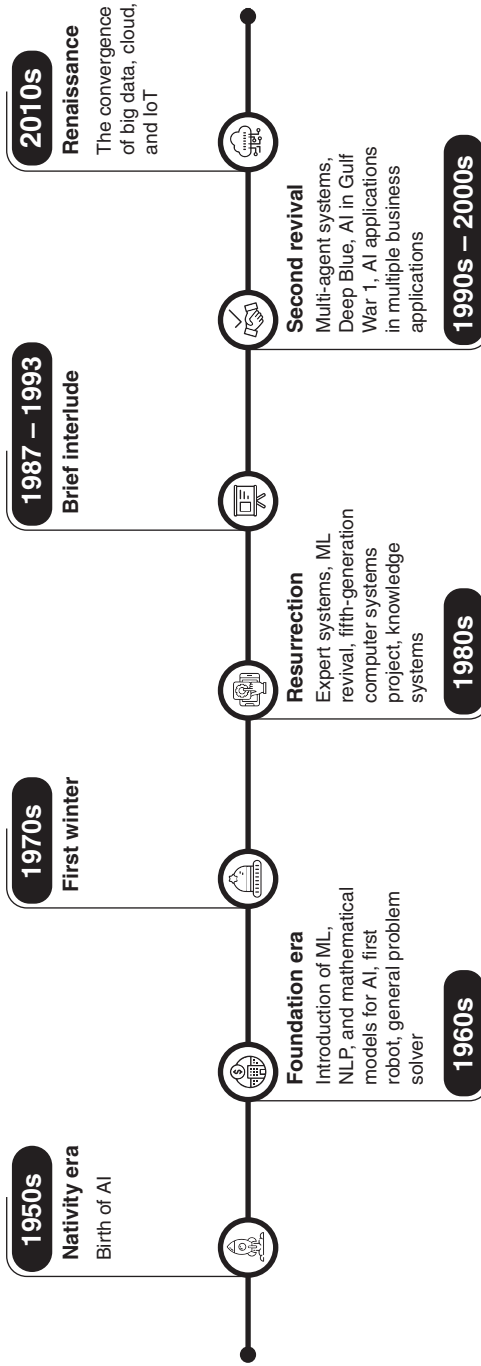


FIGURE 1.1 Evolution of AI

AI AND ITS BRANCHES

Many researchers classify AI techniques as:

- a. **Weak AI or Narrow AI.** These are applications directed to a specific problem; currently most of the practical AI implementations in the world fall under this category.
- b. **Strong AI or Artificial General Intelligence (AGI).** This is when the machine can interpret the problem just like humans and act on it. Popularly this type of transition into is also known as *singularity*.⁽⁴⁾ This is still in the lab environment and may take several decades before any meaningful application comes to life. There are many ethical issues involved with AGI, too.
- c. **Superintelligence.** This is when the machines overtake the capacity and capabilities of the human mind, going further beyond the state of singularity; this is still in early research stages.

AI deals with a variety of problem types. There are different methods to solve these different types of problems. The differences in methods lead to the different branches or disciplines within AI. Let us start with the problem types at the topmost level and their corresponding methods (hence branches) of AI:

- a. Predict a future state based on past data – **machine learning**
- b. Process human language either as text and/or voice – **natural language processing**
- c. Understand and interpret images – **image perception**
- d. Infer insights – **reasoning systems**
- e. Organize knowledge and use it to solve complex problems – **knowledge-based systems**
- f. Mimic human-like capabilities using machines – **robotics**
- g. Plan and execute complex tasks in an unknown environment – **planning**

These AI methods are not uniquely distinctive and are often used in conjunction to solve problems. However, there is enough specificity for them to qualify as separate branches of AI.

The discipline of machine learning has been extensively talked about. Within machine learning, there are many approaches like:

- **Supervised learning.** This technique is applied when there are very specific well-labeled datasets available for the program to learn the pattern and predict the output. Regression, decision tree, nearest neighbor, and support vector machines are some of the popular supervised learning algorithms.
- **Unsupervised learning.** This technique is applied when data is not properly labeled and very little human input is available but the program has to find the most probable pattern and outcome. Clustering, anomaly detection, and neural network algorithms are frequently employed for unsupervised learning.
- **Reinforcement learning.** This technique is applied when there is quite a bit of variability and ambiguity around the data, a lot more than for using supervised and unsupervised learning. In this method, the program takes feedback intermittently during the execution of the program and adapts its logic. This technique is the most popular one for designing multi-agent systems and applications like autonomous driving. Most AI techniques with the

data scientist start with a model; in reinforcement learning, one does not have to start with a model, and the model can emerge during the analysis.

The diagram in Figure 1.2 is a handy depiction of the various branches of AI.

For the last several years, deep learning has been very talked-about in the AI scene. It is one of the more popular reinforcement learning techniques today. Companies like DeepMind have contributed a lot to the popularity of deep learning. Deep learning uses artificial neural networks and convolutional neural networks. Deep learning is inspired by how biological systems learn. It works in layers and through multiple passes of the learning algorithms over the data. With each layer of analysis and learning, there is a higher level of abstraction of the raw input data to create new insights. This way the program can find correlations between actions by interpreting the corresponding data. This approach allows deep learning to work without predefined constraints and find new boundary conditions of optimal performance or achieving an objective. In rule-based systems, we create the rules based on history or understanding; deep learning helps us explore the unknown possibilities. For example, in the famous example of DeepMind optimizing HVAC systems using deep learning, they were able to find new ways of substantially reducing energy consumption by changing the sequence of operations and flow of chilled water using rules that were identified by the AI programs instead of being preset by humans. Deep learning is emerging as one of the more popular AI techniques in tackling complex problems that are not fully understood. There are many different types of deep learning algorithms like multilayer perceptron, convolutional neural networks, recurrent neural networks, autoencoder, long-short-term memory, deep belief networks, etc. Deep learning techniques have also been employed to enrich supervised and unsupervised learning in addition to being one of the most popular reinforcement learning techniques.

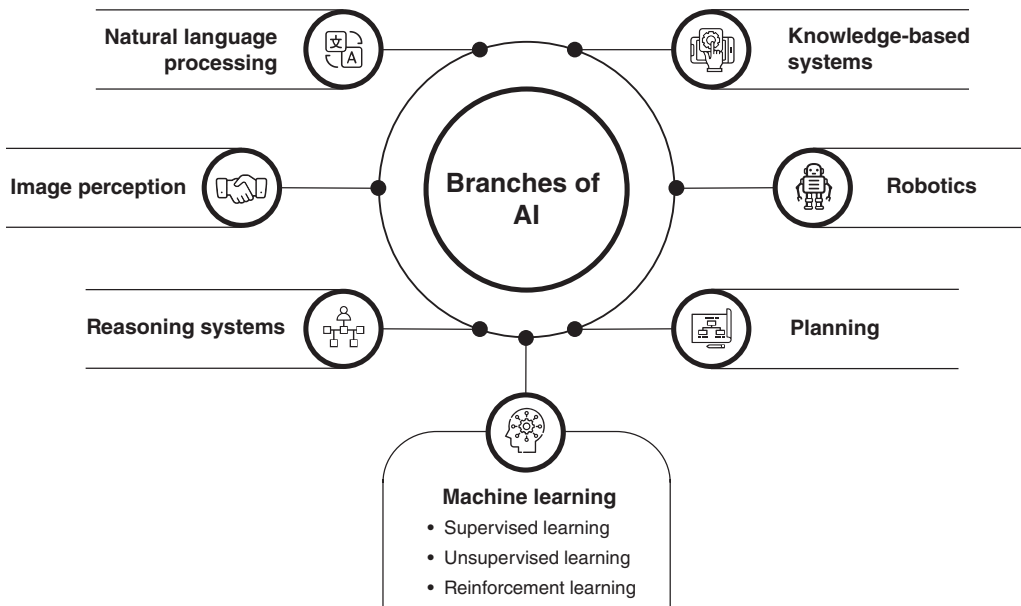


FIGURE 1.2 Branches of AI



A BIT ABOUT ALGORITHMS

Let us now clarify a few things about algorithms. The myriad techniques and algorithms can get very confusing for the uninitiated. Simply put, algorithms are the mathematical expressions that capture the process and logic chain for decision making. AI techniques, like the ones described above, take those algorithms, use data, and create insights from the decision-making process. Algorithms are the tools while data provides the fuel for the vehicle of AI to function.

Here are the top-nine things to keep in mind about algorithms:

1. **The same algorithm can be used in multiple AI techniques.** For example, the Convolutional Neural Network (CNN) algorithm can be used for both classifying and recognizing images as well as predicting the future value of something. So, the same algorithm can be applied differently for finding out which is the picture of a cat, a dog, and a traffic light, and at the same time predict how much energy your home will most likely consume tomorrow.
2. **Multiple algorithms can be applied to solve the same problem class.** For example, two different algorithms, like random forest and support vector machines, have been applied successfully to predict equipment failure. However, rarely are multiple algorithms applied to the same problem set simultaneously.
3. **Certain algorithms are better suited for specific problems.** For example, the CNN algorithm from point 1 above is better suited for energy prediction modeling as opposed to the image recognition problem.
4. **Each algorithm has a general method and a specific architecture as it gets applied to solve specific problems.** Every algorithm has a standard mathematical analysis process. However, during that analysis, it can use different variables and computations in different sequences, changing its architecture.
5. **The same algorithm can have different architectures.** Therefore it can be adapted to be applied differently for different problem classes using different techniques.
6. **Effectiveness in selecting the best algorithm for a problem comes from experience.** Data scientists with years of experience and exposure to a variety of problems can intuitively figure out the best approach. Some tools can help you navigate that, but most of them are still in their infancy.
7. **Algorithmic development is an iterative process.** Even the most experienced data scientists will refrain from claiming success at the first pass. They will experiment with multiple techniques, multiple algorithms, different architectures, and different cuts of the dataset before they finally work on a solution. This is also a very effective way of approaching this space because there are multiple choices.
8. **The algorithmic process starts with feature engineering.** This is the process of determining which variables among all available in a dataset have more relevance and better-quality data for the analysis and desired outcome. Some of the algorithms, like the deep learning algorithm LSTM, self-determine which variables are more important and relevant for the algorithm architecture.
9. **Usually, the quality of data from variables determined at the feature engineering stage is more important for most algorithms and most techniques.** However, for

certain techniques like deep learning, getting access to large volumes of data collected over time is equally important. Compared to classical algorithms like decision-tree or regression models that require sample data in tens of thousands, LSTM-type algorithms used in deep learning need samples to the tune of tens of millions.

Unless you have a strong background in statistics, understanding algorithms and the techniques can be challenging, and that is why we have data scientists. But it is important for every business and technology leader to have a basic grounding in the subject to engage intelligently with their data sciences teams.

CRITICAL SUCCESS FACTORS FOR AI INITIATIVES

Even though the discipline of AI is more than 70 years old, its broad application in the business world is more recent. There are hundreds of examples of successful AI projects, but there are precious few that are enterprise-wide and have made a big impact on a company. We do see a lot of great success stories in the government and smart city space, especially with surveillance. For example, China has one of the most incredible surveillance networks with hundreds of millions of cameras across the country and can not only monitor people but also predict potential problems. Similarly, there are many successful case studies in space research, defense, and scientific exploration.

Here are some examples from the business world where AI is at the heart of the present and future success of the enterprise:

- **Uber.** All routing, pricing, and time estimation functions in the app are now made by AI.
- **Amazon.** The entire shopping experience, dynamic pricing decisions, and recommendation engine that drives the massive amount of cross-selling are based on AI. Amazon.com is probably the biggest power user of AI in the world.
- **Tesla.** The majority of the driving features of Tesla today are based on AI and the future self-driving cars from Tesla will be fully run by AI. They are one of the most sophisticated users of computer vision, robotics, and machine learning.
- **Facebook.** All the recommendations on Facebook, whether it be for new potential friends or people to follow or ads for products and services or travel recommendations or even posts, are all determined by AI.
- **JPMC.** The bank has completely changed its customer service function and many of its internal operational processes by using AI techniques. We will talk more about this in a subsequent chapter on how AI will transform financial services.

In the last few years, the phrase *digital transformation* has been thrown around a lot. Now, this is being used more frequently with AI. We talk about that a bit in this book, too. While there is no universally accepted definition of digital transformation, here we define it as the integration of products, processes, and strategies within an organization by leveraging digital technologies such as the cloud, the edge, IoT, digital twinning, AI, and more. When AI is the backbone to make businesses more intelligent by way of accessing technology to make smart decisions, we call it AI-led digital transformation. This includes creating a data-enabled environment and

analyzing captured data to make meaningful predictions and choices. The Uber example from above is a perfect example of such a transformation.

We have studied hundreds of projects and scores of companies implementing AI. Through this exercise, we have found some best practices that lead to a higher probability of success with the AI initiatives, and missing them usually leads to disappointing results. Here are some of the key ones:

- a. The problem should be big but well-defined.
- b. First figure out the business value of solving the problem before you unleash the data scientists.
- c. Choosing the right algorithm matters a lot.
- d. Involve partners, but keep control over the data and algorithms.
- e. Change management is critical.

As we go through the different chapters in the book, we will be talking more about each of these best practices.

One thing that we have learned through the evolution of AI is that while science and math have always been extremely competent, other impediments limit the impact of AI. So, it is important to have realistic expectations and be thoughtful about the implementation.

PURPOSE AND STRUCTURE OF THE BOOK

Most organizations are making AI a part of their core agenda for their survival and success in the future. But sometimes their leaders struggle with developing a roadmap for how to reimagine their world with AI. This book is intended to help both business and technology leaders navigate the complex world of AI with its myriad dimensions in several ways:

- Understand the opportunity landscape.
- Develop a framework for business transformation.
- Investigate the possibilities across multiple industries.
- Build a technology foundation and the enabling ecosystems.
- Address broader societal concerns around ethics, privacy, and security.
- Manage the change.

This book has two types of chapters – one focused on how AI will transform specific industries and the other type focused on deployment aspects of AI across many industries.

After this introductory chapter, in Chapter 2 we will discuss how to build a framework for applying AI to any business domain.

Then we will dedicate the next eight chapters to how AI will transform eight specific industry verticals:

1. Healthcare
2. Education

3. Transportation
4. Retail
5. Financial services
6. Built environment
7. Smart cities
8. Smart government and citizen services

We chose these eight because they impact our individual lives and businesses more than others and are also the biggest contributors to a nation's GDP and future prosperity. We started this book before the COVID-19 crisis hit us, but as we kept writing we realized that the industries we had selected are also the ones that are most relevant in this current scenario.

Chapters 11 and 12 are the only technology-oriented chapters in this book, but they are not very deep into the technology or algorithms. Chapter 11 focuses on how to build a technology landscape while Chapter 12 goes into the nuances of ethics, privacy, and cybersecurity. Our goal is not to make you a technology expert; it is to give you enough orientation and tools for a holistic perspective.

Chapters 13, 14, and 15 focus on helping you get organized around your AI initiatives and manage change. Change management is an extremely critical part of the success of any AI initiative.

We believe there are a few unique aspects of this book:

- This is possibly the first and only book that tackles the topic of AI from a holistic business transformation perspective and focuses on multiple industries.
- This is a how-to book that contains some fundamental research and thinking into the various topics influencing and being impacted by AI.
- This book is peppered with a lot of practical examples and innovations as observed in this evolving landscape.
- This book will demystify the technical intricacies involved with AI for the readers and show them a path to success.
- This book requires rudimentary technical understanding and basic management orientation to navigate the space of AI effectively.

As you will see, there are several areas of overlap between the different industries and technology choices, which further proves the point that AI will morph the way we live and work.

We hope you share our excitement and find this journey with us to be useful.

REFERENCES

- (1) *Modeling the Impact of AI on the World Economy*; McKinsey Global Institute; September 2018.
- (2) *Dawn of The Age of Algorithms*; Sudhi Ranjan Sinha; January 2018.
- (3) *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*; John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon; August 1955.
- (4) *The Coming Technological Singularity: How to Survive in the Post-Human Era*; Verner Vinge, Winter 1993.