IN THIS CHAPTER

- Seeing the dream; getting beyond the hype of artificial intelligence (AI)
- Comparing AI to machine learning
- Understanding the engineering portion of AI and machine learning
- Delineating where engineering ends and art begins

# Chapter 1
# Getting the Real Story about AI



*Artificial Intelligence (AI)*, the appearance of intelligence in machines, is a huge topic today, and it's getting bigger all the time thanks to the success of new technologies (see some current examples at `https://thinkml.ai/top-5-ai-achievements-of-2019/`). However, most people are looking for everyday applications, such as talking to their smartphone. Talking to your smartphone is both fun and helpful to find out things like the location of the best sushi restaurant in town or to discover how to get to the concert hall. As you talk to your smartphone, it learns more about the way you talk and makes fewer mistakes in understanding your requests. The capability of your smartphone to learn and interpret your particular way of speaking is an example of an AI, and part of the technology used to make it happen is *machine learning*, the use of various techniques to allow algorithms to work better based on experience.

You likely make limited use of machine learning and AI all over the place today without really thinking about it. For example, the capability to speak to devices and have them actually do what you intend is an example of machine learning at work. Likewise, recommender systems, such as those found on Amazon, help you make purchases based on criteria such as previous product purchases or products that complement a current choice. The use of both AI and machine learning will only increase with time.

In this chapter, you delve into AI and discover what it means from several perspectives, including how it affects you as a consumer and as a scientist or engineer. You also discover that AI doesn't equal machine learning, even though the media often confuse the two. Machine learning is definitely different from AI, even though the two are related.

# Moving beyond the Hype

As any technology becomes bigger, so does the hype, and AI certainly has a lot of hype surrounding it. For one thing, some people have decided to engage in fear mongering rather than science. Killer robots, such as those found in the film *The Terminator*, really aren't going to be the next big thing. Your first real experience with an android AI is more likely to be in the form a health care assistant (`https://www.robotics.org/blog-article.cfm/The-Future-of-Elder-Care-is-Service-Robots/262`) or possibly as a coworker (`https://www.computerworld.com/article/2990849/meet-the-virtual-woman-who-may-take-your-job.html`). The reality is that you interact with AI and machine learning in far more mundane ways already. Part of the reason you need to read this chapter is to get past the hype and discover what AI can do for you today.

REMEMBER

You may also have heard machine learning and AI used interchangeably. AI includes machine learning, but machine learning doesn't fully define AI. This chapter helps you understand the relationship between machine learning and AI so that you can better understand how this book helps you move into a technology that used to appear only within the confines of science fiction novels.

Machine learning and AI both have strong engineering components. That is, you can quantify both technologies precisely based on *theory* (substantiated and tested explanations) rather than simply *hypothesis* (a suggested explanation for a phenomenon). In addition, both have strong science components, through which people test concepts and create new ideas of how expressing the thought process might be possible. Finally, machine learning also has an artistic component, and this is where a talented scientist can excel. In some cases, AI and machine learning both seemingly defy logic, and only the true artist can make them work as expected.

### YES, FULLY AUTONOMOUS WEAPONS EXIST

Before people send us their latest dissertations about fully autonomous weapons, yes, some benighted souls are working on such technologies. You'll find some discussions of the ethics of AI in this book, but for the most part, the book focuses on positive, helpful uses of AI to aid humans, rather than kill them, because most AI research reflects these uses. You can find articles on the pros and cons of AI online, such as the Towards Data Science article at `https://towardsdatascience.com/advantages-and-disadvantages-of-artificial-intelligence-182a5ef6588c` and the *Emerj* article at `https://emerj.com/ai-sector-overviews/autonomous-weapons-in-the-military/`.

If you really must scare yourself, you can find all sorts of sites, such as `https://www.reachingcriticalwill.org/resources/fact-sheets/critical-issues/7972-fully-autonomous-weapons`, that discuss the issue of fully autonomous weapons in some depth. Sites such as Campaign to Stop Killer Robots (`https://www.stopkillerrobots.org/`) can also fill in some details for you. We do encourage you to sign the letter banning autonomous weapons at `https://futureoflife.org/open-letter-autonomous-weapons/` — there truly is no need for them.

However, it's important to remember that bans against space-based, chemical, and certain laser weapons all exist. Countries recognize that these weapons don't solve anything. Countries will also likely ban fully autonomous weapons simply because the citizenry won't stand for killer robots. The bottom line is that the focus of this book is on helping you understand machine learning in a positive light.

## Dreaming of Electric Sheep

*Androids* (a specialized kind of robot that looks and acts like a human, such as Data in *Star Trek: The Next Generation*) and some types of *humanoid robots* (a kind of robot that has human characteristics but is easily distinguished from a human, such as C-3PO in *Star Wars*) have become the poster children for AI (see the dancing robots at `https://www.youtube.com/watch?v=1TckiTBaWkw`). They present computers in a form that people can *anthropomorphize* (give human characteristics to, even though they aren't human). In fact, it's entirely possible that one day you won't be able to distinguish between human and artificial life with ease. Science fiction authors, such as Philip K. Dick, have long predicted such an occurrence, and it seems all too possible today. The story "Do Androids Dream of Electric Sheep?" discusses the whole concept of more real than real. The idea appears as part of the plot in the movie *Blade Runner* (`https://www.warnerbros.com/movies/blade-runner`). However, some uses of robots today are just plain fun, as in the

Robot Restaurant show at https://www.youtube.com/watch?v=l1vvTtz8hpg. The sections that follow help you understand how close technology currently gets to the ideals presented by science fiction authors and the movies.


TECHNICAL STUFF

The current state of the art is lifelike, but you can easily tell that you're talking to an android. Viewing videos online can help you understand that androids that are indistinguishable from humans are nowhere near any sort of reality today. Check out the Japanese robots at https://www.youtube.com/watch?v=LyyytwT-BMk and https://www.cnbc.com/2019/10/31/human-like-androids-have-entered-the-workplace-and-may-take-your-job.html. One of the more lifelike examples is Erica (https://www.youtube.com/watch?v=oRlwvLubFxg), who is set to appear in a science fiction film. Her story appears on *HuffPost* at https://www.huffpost.com/entry/erica-japanese-robot-science-fiction-film_n_5ef6523dc5b6acab284181c3. The point is, technology is just starting to get to the point where people may eventually be able to create lifelike robots and androids, but they don't exist today.

## Understanding the history of AI and machine learning

There is a reason, other than anthropomorphization, that humans see the ultimate AI as one that is contained within some type of android. Ever since the ancient Greeks, humans have discussed the possibility of placing a mind inside a mechanical body. One such myth is that of a mechanical man called Talos (http://www.ancient-wisdom.com/greekautomata.htm). The fact that the ancient Greeks had complex mechanical devices, only one of which still exists (read about the Antikythera mechanism at http://www.ancient-wisdom.com/antikythera.htm), makes it quite likely that their dreams were built on more than just fantasy. Throughout the centuries, people have discussed mechanical persons capable of thought (such as Rabbi Judah Loew's Golem, https://www.nytimes.com/2009/05/11/world/europe/11golem.html).

AI is built on the hypothesis that mechanizing thought is possible. During the first millennium, Greek, Indian, and Chinese philosophers all worked on ways to perform this task. As early as the seventeenth century, Gottfried Leibniz, Thomas Hobbes, and René Descartes discussed the potential for rationalizing all thought as simply math symbols. Of course, the complexity of the problem eluded them (and still eludes us today, despite the advances you read about in Part 3 of this book). The point is that the vision for AI has been around for an incredibly long time, but the implementation of AI is relatively new.

The true birth of AI as we know it today began with Alan Turing's publication of "Computing Machinery and Intelligence" in 1950 (https://www.csee.umbc.

edu/courses/471/papers/turing.pdf). In this paper, Turing explored the idea of how to determine whether machines can think. Of course, this paper led to the Imitation Game involving three players. Player A is a computer and Player B is a human. Each must convince Player C (a human who can't see either Player A or Player B) that they are human. If Player C can't determine who is human and who isn't on a consistent basis, the computer wins.

A continuing problem with AI is too much optimism. The problem that scientists are trying to solve with AI is incredibly complex. However, the early optimism of the 1950s and 1960s led scientists to believe that the world would produce intelligent machines in as little as 20 years. After all, machines were doing all sorts of amazing things, such as playing complex games. AI currently has its greatest success in areas such as logistics, data mining, and medical diagnosis.

## Exploring what machine learning can do for AI

Machine learning relies on algorithms to analyze huge datasets. Currently, machine learning can't provide the sort of AI that the movies present. Even the best algorithms can't think, feel, present any form of self-awareness, or exercise free will. What machine learning can do is perform predictive analytics far faster than any human can. As a result, machine learning can help humans work more efficiently. The current state of AI, then, is one of performing analysis, but humans must still consider the implications of that analysis — making the required moral and ethical decisions. The "Considering the Relationship between AI and Machine Learning" section of this chapter delves more deeply into precisely how machine learning contributes to AI as a whole. The essence of the matter is that machine learning provides just the learning part of AI, and that part is nowhere near ready to create an AI of the sort you see in films.

REMEMBER

The main point of confusion between learning and intelligence is that people assume that simply because a machine gets better at its job (learning) it's also aware (intelligence). Nothing supports this view of machine learning. The same phenomenon occurs when people assume that a computer is purposely causing problems for them. The computer can't assign emotions and therefore acts only upon the input provided and the instruction contained within an application to process that input. A true AI will eventually occur when computers can finally emulate the clever combination used by nature:

- **Genetics:** Slow learning from one generation to the next
- **Teaching:** Fast learning from organized sources
- **Exploration:** Spontaneous learning through media and interactions with others

## Considering the goals of machine learning

At present, AI is based on machine learning, and machine learning is essentially different from statistics. Yes, machine learning has a statistical basis, but it makes some different assumptions than statistics do because the goals are different. Table 1-1 lists some features to consider when comparing AI and machine learning to statistics.

**TABLE 1-1:** Comparing Machine Learning to Statistics

| Technique | Machine Learning | Statistics |
|---|---|---|
| Data handling | Works with big data in the form of networks and graphs; raw data from sensors or the web text is split into training and test data. | Models are used to create predictive power on small samples. |
| Data input | The data is sampled, randomized, and transformed to maximize accuracy scoring in the prediction of out-of-sample (or completely new) examples. | Parameters interpret real-world phenomena and provide a stress on magnitude. |
| Result | Probability is taken into account for comparing what could be the best guess or decision. | The output captures the variability and uncertainty of parameters. |
| Assumptions | The scientist learns from the data. | The scientist assumes a certain output and tries to prove it. |
| Distribution | The distribution is unknown or ignored before learning from data. | The scientist assumes a well-defined distribution. |
| Fitting | The scientist creates a best fit, but generalizable, model. | The result is fit to the present data distribution. |

## Defining machine learning limits based on hardware

Huge datasets require huge amounts of memory. Unfortunately, the requirements don't end there. When you have huge amounts of data and memory, you must also have processors with multiple cores and high speeds. One of the problems that scientists are striving to solve is how to use existing hardware more efficiently. In some cases, waiting for days to obtain a result to a machine learning problem simply isn't possible. The scientists who want to know the answer need it quickly, even if the result isn't quite right. With this in mind, investments in better hardware also require investments in better science. This book considers some of the following issues as part of making your machine learning experience better:

- **Obtaining a useful result:** As you work through the book, you discover that you need to obtain a useful result first, before you can refine it. In addition, sometimes tuning an algorithm goes too far and the result becomes quite fragile (and possibly useless outside a specific dataset).
- **Asking the right question:** Many people get frustrated in trying to obtain an answer from machine learning because they keep tuning their algorithm without asking a different question. To use hardware efficiently, sometimes you must step back and review the question you're asking. The question might be wrong, which means that even the best hardware will never find the answer.
- **Relying on intuition too heavily:** All machine learning questions begin as a hypothesis. A scientist uses intuition to create a starting point for discovering the answer to a question. Failure is more common than success when working through a machine learning experience. Your intuition adds the art to the machine learning experience, but sometimes intuition is wrong and you have to revisit your assumptions.


When you begin to realize the importance of environment to machine learning, you can also begin to understand the need for the right hardware and in the right balance to obtain a desired result. The current state-of-the-art systems actually rely on Graphical Processing Units (GPUs) to perform machine learning tasks. Relying on GPUs does speed the machine learning process considerably. A full discussion of using GPUs is outside the scope of this book, but you can read more about the topic at `https://devblogs.nvidia.com/parallelforall/bidmach-machine-learning-limit-gpus/` and `https://towardsdatascience.com/what-is-a-gpu-and-do-you-need-one-in-deep-learning-718b9597aa0d`.

# Overcoming AI Fantasies

As with many other technologies, AI and machine learning both have their fantasy or fad uses. For example, some people are using machine learning to create Picasso-like art from photos using products like NightCafé (`https://creator.nightcafe.studio/`), which supports people who really enjoy this art form. You can read all about using machine learning to create art at `https://www.washingtonpost.com/news/innovations/wp/2015/08/31/this-algorithm-can-create-a-new-van-gogh-or-picasso-in-just-an-hour/`. Of course, the problems with such use are many. For one thing, most people wouldn't really want a Picasso created in this manner except as a fad item (because no one had done it before). The point of art isn't in creating an interesting interpretation of a particular real-world representation, but rather in seeing how the artist interpreted it. The end of the article points out that the computer can only copy an

existing style at this stage — not create an entirely new style of its own. The following sections discuss AI and machine learning fantasies of various sorts.

## Discovering the fad uses of AI and machine learning

AI is entering an era of innovation that you used to read about only in science fiction. It can be hard to determine whether a particular AI use is real or simply the dream child of a determined scientist. For example, *The Six Million Dollar Man* (https://en.wikipedia.org/wiki/The_Six_Million_Dollar_Man) is a television series that looked fanciful at one time. When it was introduced, no one actually thought that we'd have real-world bionics at some point. However, Hugh Herr (https://www.smithsonianmag.com/innovation/future-robotic-legs-180953040/) and others (https://www.fiercebiotech.com/medtech/using-onboard-ai-to-power-quicker-more-complex-prosthetic-hands) have other ideas — bionic legs and arms really are possible now. Of course, they aren't available for everyone yet; the technology is only now becoming useful. Muddying the waters is *The Six Billion Dollar Man* movie, based partly on *The Six Million Dollar Man* television series (https://www.cinemablend.com/new/Mark-Wahlberg-Six-Billion-Dollar-Man-Just-Made-Big-Change-91947.html), which has suffered delays for various reasons (https://screenrant.com/mark-wahlberg-six-billion-dollar-man-delays-updates/). The fact is that AI and machine learning will both present opportunities to create some amazing technologies and that we're already at the stage of creating those technologies, but you still need to take what you hear with a huge grain of salt.

One of the more interesting uses of machine learning for entertainment purposes is the movie *B* (https://www.cinemablend.com/news/2548939/one-sci-fi-movie-will-be-able-to-film-during-the-pandemic-thanks-to-casting-an-ai-robot-as-its-lead), which stars an android named Erica. The inventors of Erica, Hiroshi Ishiguro and Kohei Ogawa, have spent a great deal of time trying to make her lifelike by trying to implement the human qualities of intent and desire (https://www.yoichimatsuyama.com/conversation-with-evolving-robotic-species-interview-with-hiroshi-ishiguro/). The result is something that encroaches on the uncanny valley (https://www.scientificamerican.com/article/why-uncanny-valley-human-look-alikes-put-us-on-edge/) in a new way. The plot of this movie will be on the same order as *Ex Machina* (https://www.indiewire.com/2020/06/ex-machina-real-robot-erica-science-fiction-movie-1234569484/).

REMEMBER

To make the future uses of AI and machine learning match the concepts that science fiction has presented over the years, real-world programmers, data scientists, and other stakeholders need to create tools. Nothing happens by magic, even though it may look like magic when you don't know what's happening behind the scenes. In order for the fad uses for AI and machine learning to become real-world uses, developers, data scientists, and others need to continue building real-world tools that may be hard to imagine at this point.

## Considering the true uses of AI and machine learning

You find AI and machine learning used in a great many applications today. The only problem is that the technology works so well that you don't know that it even exists. In fact, you might be surprised to find that many devices in your home already make use of both technologies. Both technologies definitely appear in your car and most especially in the workplace. In fact, the uses for both AI and machine learning number in the millions — all safely out of sight even when they're quite dramatic in nature. Here are just a few of the ways in which you might see AI used:

- **Fraud detection:** You get a call from your credit card company asking whether you made a particular purchase. The credit card company isn't being nosy; it's simply alerting you to the fact that someone else could be making a purchase using your card. The AI embedded within the credit card company's code detected an unfamiliar spending pattern and alerted someone to it.
- **Resource scheduling:** Many organizations need to schedule the use of resources efficiently. For example, a hospital may have to determine where to put a patient based on the patient's needs, availability of skilled experts, and the amount of time the doctor expects the patient to be in the hospital.
- **Complex analysis:** Humans often need help with complex analysis because there are literally too many factors to consider. For example, the same set of symptoms could indicate more than one problem. A doctor or other expert might need help making a diagnosis in a timely manner to save a patient's life.
- **Automation:** Any form of automation can benefit from the addition of AI to handle unexpected changes or events. A problem with some types of automation today is that an unexpected event, such as an object in the wrong place, can actually cause the automation to stop. Adding AI to the automation can allow the automation to handle unexpected events and continue as if nothing happened.
- **Customer service:** The customer service line you call today may not even have a human behind it. The automation is good enough to follow scripts and use various resources to handle the vast majority of your questions. With

good voice inflection (provided by AI as well), you may not even be able to tell that you're talking with a computer.

- **Safety systems:** Many of the safety systems found in machines of various sorts today rely on AI to take over the vehicle in a time of crisis. For example, many automatic braking systems rely on AI to stop the car based on all the inputs that a vehicle can provide, such as the direction of a skid.
- **Machine efficiency:** AI can help control a machine in such a manner as to obtain maximum efficiency. The AI controls the use of resources so that the system doesn't overshoot speed or other goals. Every ounce of power is used precisely as needed to provide the desired services.

This list doesn't even begin to scratch the surface. You can find AI used in many other ways. However, it's also useful to view uses of machine learning outside the normal realm that many consider the domain of AI. Here are a few uses for machine learning that you might not associate with an AI:

- **Access control:** In many cases, access control is a yes or no proposition. An employee smartcard grants access to a resource much in the same way that people have used keys for centuries. Some locks do offer the capability to set times and dates that access is allowed, but the coarse-grained control doesn't really answer every need. By using machine learning, you can determine whether an employee should gain access to a resource based on role and need. For example, an employee can gain access to a training room when the training reflects an employee role.
- **Animal protection:** The ocean might seem large enough to allow animals and ships to cohabitate without problem. Unfortunately, many animals get hit by ships each year. A machine learning algorithm could allow ships to avoid animals by learning the sounds and characteristics of both the animal and the ship.
- **Predicting wait times:** Most people don't like waiting when they have no idea of how long the wait will be. Machine learning allows an application to determine waiting times based on staffing levels, staffing load, complexity of the problems the staff is trying to solve, availability of resources, and so on.

## Being useful; being mundane

Even though the movies make it sound like AI is going to make a huge splash, and you do sometimes see some incredible uses for AI in real life, the fact of the matter is that most uses for AI are mundane, even boring. For example, a recent article

details how Verizon uses the R language to analyze security breach data (https://www.computerworld.com/article/3001832/data-analytics/how-verizon-analyzes-security-breach-data-with-r.html and https://softwarestrategiesblog.com/category/verizons-2020-data-breach-investigations-report-dbir/). Part 5 of this book provides you with real-world examples of this same sort of analysis. The act of performing this analysis is dull when compared to other sorts of AI activities, but the benefits are that Verizon saves money performing the analysis using R, and the results are better as well.

In addition, Python developers (see Chapters 4 and 5 for Python language details) have a huge array of libraries available to make machine learning easy. In fact, Kaggle (https://www.kaggle.com/competitions) provides competitions to allow Python developers and R practitioners to hone their machine learning skills in creating practical applications. The results of these competitions often appear later as part of products that people actually use. Although R still relies on strong support from the statistical community in academic research, the Python development community is particularly busy creating new libraries to make development of complex data science and machine learning applications easier (see https://www.globalsqa.com/top-20-open-source-python-libraries/ for the top 20 Python libraries in use today).

# Considering the Relationship between AI and Machine Learning

Machine learning is only part of what a system requires to become an AI. The machine learning portion of the picture enables an AI to perform these tasks:

- Adapt to new circumstances that the original developer didn't envision
- Detect patterns in all sorts of data sources
- Create new behaviors based on the recognized patterns
- Make decisions based on the success or failure of these behaviors

The use of algorithms to manipulate data is the centerpiece of machine learning. To prove successful, a machine learning session must use an appropriate algorithm to achieve a desired result. In addition, the data must lend itself to analysis using the desired algorithm, or it requires a careful preparation by scientists.

AI encompasses many other disciplines to simulate the thought process successfully. In addition to machine learning, AI normally includes

- **Natural language processing:** The act of allowing language input and putting it into a form that a computer can use.
- **Natural language understanding:** The act of deciphering the language in order to act upon the meaning it provides.
- **Knowledge representation:** The ability to store information in a form that makes fast access possible.
- **Planning (in the form of goal seeking):** The ability to use stored information to draw conclusions in *near real time* (almost at the moment it happens, but with a slight delay, sometimes so short that a human won't notice, but the computer can).
- **Robotics:** The ability to act upon requests from a user in some physical form.

In fact, you might be surprised to find that the number of disciplines required to create an AI is huge. Consequently, this book exposes you to only a portion of what an AI contains. However, even the machine learning portion of the picture can become complex because understanding the world through the data inputs that a computer receives is a complex task. Just think about all the decisions that you constantly make without thinking about them. For example, just the concept of seeing something and knowing whether you can interact successfully with it can become a complex task.

# Considering AI and Machine Learning Specifications

As scientists continue to work with a technology and turn hypotheses into theories, the technology becomes related more to *engineering* (where theories are implemented) than *science* (where theories are created). As the rules governing a technology become clearer, groups of experts work together to define these rules in written form. The result is *specifications* (a group of rules that everyone agrees upon).

Eventually, implementations of the specifications become *standards* that a governing body, such as the IEEE (Institute of Electrical and Electronics Engineers) or a combination of the ISO/IEC (International Organization for Standardization/International Electrotechnical Commission), manages. AI and machine learning

have both been around long enough to create specifications, but you currently won't find any standards for either technology. However, you can find plans for such standards in places like National Institute of Standards and Technology (NIST) at https://www.nist.gov/topics/artificial-intelligence/ai-standards.

The basis for machine learning is math. Algorithms determine how to interpret big data in specific ways. The math basics for machine learning appear in Part 3 of this book. You discover that algorithms process input data in specific ways and create predictable outputs based on the data patterns. What isn't predictable is the data itself. The reason you need AI and machine learning is to decipher the data in such a manner to be able to see the patterns in it and make sense of them.

You see the specifications detailed in Part 4 in the form of algorithms used to perform specific tasks. When you get to Part 5, you begin to see the reason that everyone agrees to specific sets of rules governing the use of algorithms to perform tasks. The point is to use an algorithm that will best suit the data you have in hand to achieve the specific goals you've created. Professionals implement algorithms using languages that work best for the task. Machine learning relies on Python and R, and to some extent MATLAB, Java, Julia, and C++. (See the discussion at https://www.quora.com/What-is-the-best-language-to-use-while-learning-machine-learning-for-the-first-time for details.)

# Defining the Divide between Art and Engineering

The reason that AI and machine learning are both sciences and not engineering disciplines is that both require some level of art to achieve good results. The artistic element of machine learning takes many forms. For example, you must consider how the data is used. Some data acts as a baseline that trains an algorithm to achieve specific results. The remaining data provides the output used to understand the underlying patterns. No specific rules governing the balancing of data exist; the scientists working with the data must discover whether a specific balance produces optimal output.

REMEMBER

Cleaning the data also lends a certain amount of artistic quality to the result. The manner in which a scientist prepares the data for use is important. Some tasks, such as removing duplicate records, occur regularly. However, a scientist may also choose to filter the data in some ways or look at only a subset of the data. As a result, the cleaned dataset used by one scientist for machine learning tasks may not precisely match the cleaned dataset used by another.

You can also tune the algorithms in certain ways or refine how the algorithm works. Again, the idea is to create output that truly exposes the desired patterns so that you can make sense of the data. For example, when viewing a picture, a robot may have to determine which elements of the picture it can interact with and which elements it can't. The answer to that question is important if the robot must avoid some elements to keep on track or to achieve specific goals.

When working in a machine learning environment, you also have the problem of input data to consider. For example, the microphone found in one smartphone won't produce precisely the same input data that a microphone in another smartphone will. The characteristics of the microphones differ, yet the result of interpreting the vocal commands provided by the user must remain the same. Likewise, environmental noise changes the input quality of the vocal command, and the smartphone can experience certain forms of electromagnetic interference. Clearly, the variables that a designer faces when creating a machine learning environment are both large and complex.

The art behind the engineering is an essential part of machine learning. The experience that a scientist gains in working through data problems is essential because it provides the means for the scientist to add values that make the algorithm work better. A finely tuned algorithm can make the difference between a robot successfully threading a path through obstacles and hitting every one of them.

# Predicting the Next AI Winter

Development of machine learning and AI is slow for a number of reasons, such as a lack of powerful hardware, lack of suitable data to feed algorithms, and people's inability to understand their own thought processes. Businesses, however, are looking for ways to generate cash quickly based on new technologies. Obviously, slow development doesn't work well with a quick return on investment (ROI). Developer-entrepreneurs exacerbate the problem by overselling technologies. They indicate that the state of the art is more advanced than it really is, often to enjoy windfall profits, gain power, and advance their careers. Because of the difference between timing and expectations, machine learning and AI have both experienced *AI winters*, a period of time when business shows little or no interest in the development of new processes, technologies, or strategies.

The first AI winter happened as a result of unfulfilled expectations resulting from the overselling of the technology and unanticipated difficulties. During the summer of 1956, various scientists attended a workshop held on the Dartmouth College campus to create artificially intelligent machines. They predicted that machines that could reason as effectively as humans would require, at most, a

generation to come about. They were wrong. Only now have we realized machines that can perform mathematical and logical reasoning as effectively as a human. To achieve true human understanding, an AI would also need to demonstrate intelligence in the visual-spatial, bodily-kinesthetic, creative, interpersonal, intrapersonal, and linguistic realms. The stated problem with the Dartmouth College and other endeavors of the time relates to hardware — the processing capability to perform calculations quickly enough to create a simulation. However, that's not really the whole problem. Yes, hardware does figure in to the picture, but you can't simulate processes that you don't understand, especially if you lack suitable data. Even so, the reason that AI is somewhat effective today is that the hardware has finally become powerful enough to support the required number of calculations.

REMEMBER

Anyone who has spent a lot of time analyzing the machine learning and AI fields knows that the current technology has reached a kind of plateau. The technology continues to advance incrementally, but there aren't any true new uses for either machine learning or AI right now. On the other hand, businesses are effectively using both machine learning and AI to generate a profit. So, some people feel that machine learning and AI are headed toward another AI winter because of unfulfilled expectations and overselling (think about the self-driving car), while others feel that business actually is satisfied with the progress currently being made (think about the use of recommender systems on sites such as Amazon.com).

Sites such as `https://www.thinkautomation.com/bots-and-ai/the-ai-winter-is-coming/` see an AI winter in the near future, partly because the terms *machine learning, deep learning,* and *AI* have become overused and ill-defined. These same sites look at how business is actually using machine learning and AI today. In most cases, these sources say that the technologies are used for background processes, not front-line customer interactions. The thought is that automation used for front-end processes isn't actually machine learning or AI, and that companies will eventually see automation as being separate from machine learning and AI. As a result, they will again stop investing in either technology. In many cases, proponents of an upcoming AI winter state that scientists should focus on the amazing array of tasks that machine learning and AI can perform today, rather than continue to hype some nebulous future tasks.

Before you get the idea that everyone is expecting another AI winter, you need to look at the other side of the argument. Sites such as `https://towardsdatascience.com/there-wont-be-an-ai-winter-this-time-332a4b6d6f07` are saying that machine learning and AI are both so deeply embedded that an AI winter really isn't possibly any longer. Typically, the articles you see are forthright in stating that machine learning and AI haven't met certain goals, like creating autonomous vehicles. Even though these goals aren't feasible today, the potential exists for achieving them in the future when scientists have completed more research.

Moreover, because of the research conducted and the applications created, both machine learning and AI have become profitable, so business will continue to support them. The Towards Data Science article is good because it points out a wealth of vendors who are actually using machine learning in major line-of-business applications that generate huge profits.

In thinking about the future of machine learning and AI, considering a more moderate approach is likely best, such as the one found at `https://mindmatters.ai/2020/01/so-is-an-ai-winter-really-coming-this-time/`. At this point, data scientists and other researchers need to take a step back and consider the next level. The current technologies can only take us so far. They're profitable, but they can't produce a self-driving car and they certainly can't produce a robot of the intelligence found in the film *Ex Machina.* So, if there is an AI winter, it's likely to be a mild one because companies like Amazon.com and Google aren't going to throw their technologies out because a few reporters think that they should. In short, the concepts, ideas, and technologies that you discover in this book remain viable and allow you to move forward in a career of your choice.