CHAPTER 1

Basics of Artificial Intelligence

Fashion not only provides functional purpose, but captures mysterious and elusive aspects of being human. Fashion expresses and invokes human emotion and creativity. How we look and sometimes even how we feel is intertwined in this industry. Fashion has always been forward looking, grabbing onto new technologies as they arise. Artificial intelligence is no exception, and it's moving as quickly as fashion does.

Artificial intelligence (**AI**) is a field of computer science that looks at the logic behind human intelligence. The field seeks ways to understand how we think and to re-create this intelligence in machines. Because of its nature, AI extends across human activities, making it relevant in different ways to every industry.

The intersection of fashion and AI is a rich and expansive space that is just beginning to be explored. As AI continues to develop, it becomes harder to comprehend for nontechnical followers. The challenge of comprehension stands in the way of meaningful developments between these two fields.

This chapter briefly covers basic concepts in artificial intelligence to provide a foundation for understanding its applications in the fashion industry. The rest of the book expands on these ideas and more.

L. Luce, *Artificial Intelligence for Fashion*,
https://doi.org/10.1007/978-1-4842-3931-5_1

Why Does AI Matter?

In “The State of Fashion 2018,” a report by McKinsey & Company and The Business of Fashion, 75% of retailers plan to invest in artificial intelligence over 2018 and 2019. It is changing the way the fashion industry does business across the entire fashion value chain. Providing customized experiences and better forecasting is just the start.

Currently, up to 30% of activities in 60% of occupations across all industries can be automated. It will still take time to implement some of this automation and reskill the current workforce. At this rate, there is no question that artificial intelligence will significantly impact the way we work.

What Is AI?

Artificial intelligence has become a confusing term. Machine learning, deep learning, and artificial intelligence are terms often used interchangeably, which may leave to question, what is the difference?

Machine learning is a way of achieving AI. In 1959, it was defined by Arthur Samuel as “the ability to learn without being explicitly programmed.” Usually this is done through “training.” **Deep learning** is an approach to machine learning, which usually involves large neural networks. Figure 1-1 shows a graphical representation of the relationship between AI, machine learning, and deep learning.

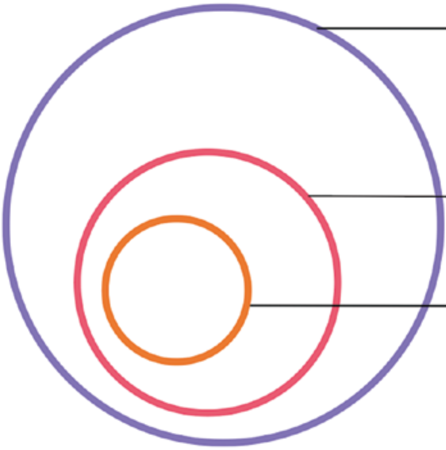

Figure 1-1. *The relationship between AI, machine learning, and deep learning*

Machine Learning

Machine learning makes up a large portion of artificial intelligence being applied in businesses today. The goals of machine learning are to automate processes in order to decrease human effort, and to discover complex patterns that humans cannot interpret on their own.

This analogy is not perfect, but you can think of it this way: machine learning is to programming as the sewing machine is to sewing. Before the advent of the sewing machine, every stitch was sewn by hand. Once the sewing machine was introduced, sewing became faster, because not every stitch was handled by a human. With machine learning, we can build programs that handle far more complexity without having to hand **code** every detail. Ultimately, however, seams can't sew themselves, and machine learning continues to require a human hand to make it work.

In machine learning, machines are used to identify patterns in data and frequently predict the values of nonexistent data, often correlating to events happening in the future. Machine learning encompasses many methods for learning from data and makes up a large portion of research happening in artificial intelligence today.

What Is Intelligence?

> *The true sign of intelligence is not knowledge but imagination.*
>
> —Albert Einstein

While we intuitively know what intelligence is, it turns out to be difficult to summarize or formally define. There are many theories and definitions about what makes humans intelligent. How to measure intelligence has been argued by philosophers for centuries.

Shane Legg and Marcus Hutter collected over 70 experts' definitions of intelligence in a paper called "A Collection of Definitions of Intelligence." In an effort to derive a single definition, they came up with this: *Intelligence measures an agent's ability to achieve goals in a wide range of environments.*

In artificial intelligence, systems are often designed to mimic behaviors of the human mind. Researchers look to the human mind as a model of intelligence. The original goal of reconstructing human intelligence in machines requires teaching machines to carry out many complex functions. Reasoning, problem solving, memory recall, planning, learning, processing natural language, perception, manipulation, social intelligence, and creativity are all part of reaching this goal.

The Turing Test

How can we know if a machine is intelligent? The **Turing test** (**TT**) was proposed by Alan Turing in 1950 as one of the first tests of intelligence in machines. It is a challenge to understand whether a machine acts like a human. To pass the test, a human interrogator asks questions to the machine. If the human interrogator cannot distinguish which responses are from a human and which are from a machine, the machine passes the test.

The Turing test has appeared time and time again in popular science-fiction movies over the past 40 years. *Ex Machina* and *Blade Runner* are examples. It is one of many "Are we there yet?" checkpoints for the field.

How Machines Learn

> *Making mental connections is our most crucial learning tool, the essence of human intelligence; to forge links; to go beyond the given; to see patterns, relationships, context.*
>
> —Marilyn Ferguson, author

Understanding human behavior is complicated because humans do not always act rationally or logically. We can improve a machine's ability to predict human behavior by searching for patterns. These patterns help to discover and define trends. By analyzing these trends and modeling them with algorithms, machines can mimic human responses to certain inputs. Then, when encountering these inputs in real-world contexts, they are able to respond accordingly.

What Is Learning?

If we could simplify human learning, we might say that humans take information from their environment, relate it to something, and then learn or act. These inputs could be something they see, smell, taste, hear, feel, or even their interpretation of a mood or tone. That information is related to prior knowledge a person has about the world, making a connection. From there, a human might act on their new knowledge, explore, or innovate. This process can be observed in Figure 1-2.

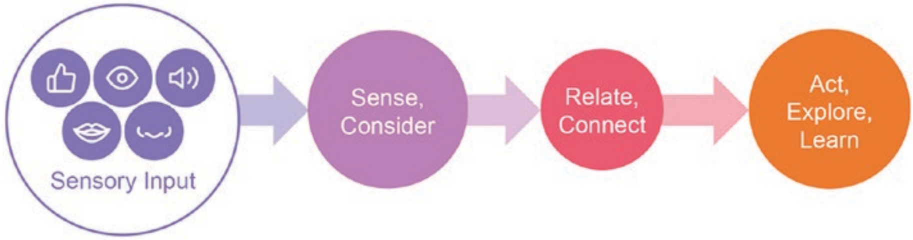

Figure 1-2. *How humans learn*

Machines are given input in the form of data. The machine interprets that data and learns from it. Then the machine evaluates that data before outputting the information that has been defined as useful to a human to interpret. This is the prediction phase, as shown in Figure 1-3.

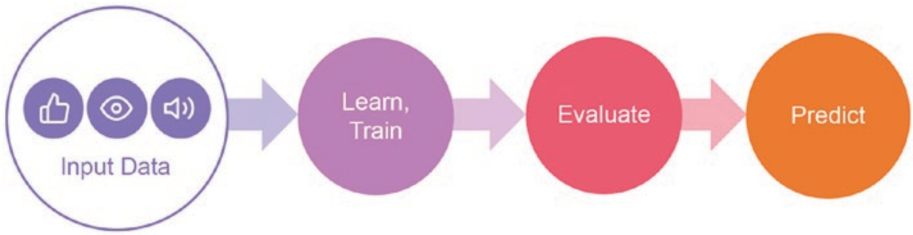

Figure 1-3. *How machines learn*

Where does the data come from? Machines are collecting data through **hardware** inputs as well as **software** programs. You can think of the hardware as the body, and software as the mind of a machine. Hardware addresses the area of **machine perception**, and software addresses the idea of both **machine language** and human language.

Machine Perception

Machines can perceive the environment through sight, feeling, and hearing via **sensors**. Sensors are a part of a machine's hardware system. They measure physical events like temperature, pressure, force, acceleration, sound, and light.

In fact, your phone can measure almost all these things. Phones sense through small electronics called **microelectromechanical systems** (**MEMS**). The microphone, camera, **inertial measurement units** (or **IMUs**, which help track position), and **proximity sensors** are all examples of MEMs. These sensors can also be found in various Internet of Things (IoT) devices.

In collaboration with these sensors, software systems on a machine can do things like interpret when a phone is upside down or right side up, measure human locomotion, and detect faces or sounds.

Language

Human languages are critical for communication. We use words and phrases, combining them in multiple ways in order to express ideas and emotions. Machines use machine languages to define models and parameters. Human language and machine perception both provide inputs in the form of data for machines to use to learn from.

An important distinction exists between machine languages and human languages. Machine languages are written in code. Originally, this code was only a series of 0s and 1s, or binary. Different combinations of 0s and 1s encode different information to machines. Over time, humans have created programming languages that interface between human language and machine language to make the job of coding easier.

The output of a machine is most useful when it can be interpreted by a human, which makes human language a useful concept for machines to understand.

Topics in Artificial Intelligence

> *The role of the computer is not to displace human creativity but rather to amplify it.*
>
> —Ray Kurzweil, *The Age of Intelligent Machines*

Successfully applying AI today requires understanding which techniques should be used for solving a given problem. There is currently no single algorithm that will provide value in every aspect of the fashion industry. The term *AI* as an overarching category can be confusing because it often leads people to believe that AI is a mysterious black box that can solve any problem. In reality, it is made up of several application areas, tools, and techniques. Understanding the broader categories and more specific subcategorization of the field gives a picture of how it all fits together.

Application areas discussed in this book include:

- Natural language processing (NLP)
- Computer vision (CV)
- Predictive analytics
- Robotics

Some commonly used tools and techniques include:

- Neural networks
- Generative adversarial networks (GANs)
- Data mining

Not every topic in the field of AI is covered in this book. Because the categories often overlap, this book is written with more information being introduced cumulatively as you read through. Later chapters may rely on explanation from earlier chapters.

Application Areas

The phrasing "application areas" refers to the specific areas where machine learning tools and techniques can be applied. Natural language processing, computer vision, predictive analytics, and robotics might use some of the same techniques like neural networks to solve different types of problems. Those application areas can be further extended to industry applications.

Natural Language Processing

Machine language and human language meet in **natural language processing** (**NLP**). NLP is a way for computers to comprehend human languages. Every day, our interactions on the Web—things we post on social media, text messages we write, and so forth—contribute to an ever-expanding mass of data. Of this, it is estimated that 80% of the 2.5 quintillion bytes of data created every day is **unstructured data.** This is written in free form, unorganized and historically hard to parse. We can use NLP to understand the content and context of this unstructured data, unlocking a rich treasure trove of information about ourselves.

Natural language processing is applied in multiple product categories including conversational shopping, AI customer service chatbots, and virtual assistants and stylists.

Computer Vision

Computer vision (**CV**) is used to process and analyze images and videos. CV automates tasks we might associate with the human visual system and more. Although computer vision is a field of its own, artificial intelligence has played a major role in recent progress. Computer vision is frequently used in fashion applications because the fashion industry is so visual.

In the fashion industry, computer vision is being used in technologies such as visual search, smart mirrors, social shopping, trend forecasting, virtual reality, and augmented reality.

Predictive Analytics

> *AI can identify upcoming trends faster than industry insiders to enhance the design process.*
>
> —Avery Baker, chief brand officer, Tommy Hilfiger

Predictive analytics uses a variety of methods that use historical information to predict events that will happen in the future. These methods range in complexity and include data mining, basic statistics, and machine learning.

In this book, predictive analytics show up in two other areas: recommender systems and demand forecasting.

Recommender systems are part of predictive analytics. They seek to understand user or customer behavior and recommend products or services that the user is likely to like or purchase. Recommender systems have played a critical role for discovering products in e-commerce. You'll find them everywhere, from fashion retail web sites to behind the scenes in subscription box services. You'll also notice them in other areas including music and video streaming on sites like Netflix or YouTube.

Demand forecasting is used to optimize supply-chain planning. By predicting demand for products, the fashion industry can reduce overproduction, thereby cutting costs and reducing waste.

Robotics

Robotics, especially in apparel manufacturing, is a unique area of study that requires domain expertise across fashion, mechanical engineering, and machine learning. Robots have been used in industrial settings for many years in manufacturing for automotive, aerospace, and other industries that deal with mostly rigid parts.

Robotic manufacturing in the fashion industry is still a nascent field because of the complexities involved with handling fabrics. Nonetheless, with improvements in computer vision and in the planning algorithms needed to perform complex tasks, robotics is being adopted in fashion.

Tools and Techniques

Tools and techniques like neural networks, generative adversarial networks, and data mining are used across application areas. These methods are constantly changing and evolving to return higher quality results in industry.

Neural Networks

Neural networks are a subcategory of machine learning**.** They were originally modeled after our understanding of the behavior of neurons in the human brain: in the brain, a single neuron takes in input, processes it, and sends output. Neuroscience has moved away from this idea. We now know brains don't actually work like this, and the statistics behind neural networks in machine learning have been developed independently of neuroscience.

Neural networks are typically created with layers that compute information in parallel. They're composed of interconnected **nodes**. Knowledge in these systems is represented by the patterns that are taken on by nodes passing information to each other.

The way people think about the composition of a neural network usually includes three basic parts:

- *Input layers*: Contain input data
- *Hidden layers*: Contain the synapse architecture
- *Output layers*: Provide results from the network

Within that framework, a neural network can take on many architectures. Not all neural networks are the same. In the implementation, **training** is also an important part of the process. Training involves sending data through the neural network. In this stage, the network is learning complex connections between inputs and desired outputs. In many instances, the network's effectiveness is reliant on high-quality data.

Neural networks are frequently used in the application areas discussed earlier in this chapter. Understanding the basic mechanisms of neural networks helps provide a foundation for understanding how contemporary artificial intelligence works.

Generative Adversarial Networks

Unsupervised learning can be inefficient because machines must learn by themselves. What is obvious to us may not be obvious to a machine. **Generative adversarial networks** (**GANs**) are one way to increase the efficiency of unsupervised learning. GANs use two neural networks: one network generates results, and the other evaluates the accuracy of those results.

GANs are a more recently adopted technology in the machine learning space and have been proposed by companies such as Amazon as a method for creating AI fashion designers in 2017. These and other generative models are especially promising for creating unique new images as well as for filling in information from images that are incomplete or damaged.

Data Mining

Data is critical to any task in machine learning. Without data, the machine has nothing to train from. Data can include information such as video, images, and text. **Data collection** refers to the process of collecting data for analysis.

In many cases, data collection is just the beginning. What do you do with all the data? **Data mining** is about uncovering useful information in large amounts of data. For the fashion industry, social media can be a treasure trove for learning about the way customers feel about products and trends.

INTRODUCING BETTY & RUTH

Throughout the book, you'll find mention of a fictitious women's fashion brand called Betty & Ruth. Through the narrative of this brand, there will be examples of how certain tasks are currently handled at the fashion company and how they might be improved using techniques discussed in the book.

The examples from Betty & Ruth are an opportunity to explore implementation and sometimes expand on how other related technologies might fit into the picture.

Summary

Different application areas, techniques, and tools have strengths and weaknesses at specific tasks. Application areas of AI are often targeted at addressing specific needs (for example, image- vs. language-related problems are addressed using computer vision and natural language processing respectively). Successfully applying AI today requires understanding which techniques and tools make sense for your application.

Computer vision is an inherently visual field of artificial intelligence and is applied in instances dealing with images and video. Natural language processing, on the other hand, deals with communicating between human languages and machine languages.

Artificial intelligence, and especially machine learning methods, use data and models to understand and make predictions about questions we don't have the answers to.

For a deeper dive in the basics of artificial intelligence, I recommend the following books. You can find even more in the annotated bibliography and in the upcoming chapters of this book.

- *Artificial Intelligence: A Modern Approach*, by Stuart Russell and Peter Norvig (Pearson, 2016)
- *Artificial Intelligence: The Basics* by Kevin Warwick (Routledge, 2011)

Terminology from This Chapter

Artificial intelligence (**AI**)—A field of computer science that aims to teach machines to behave intelligently.

Code—Languages that are interpretable by machines.

Computer vision (**CV**)—A field of computer science dealing with the visual system. This includes teaching computers to process, analyze, and understand images and videos.

Data—Information that can be measured, collected, reported, and analyzed. It can take the form of various media, including text, images, and video.

Data collection—A process of gathering data for analysis.

Data mining—Uncovering useful information within a large dataset. *Refer to Chapter 9 for more information.*

Deep learning—Machine learning methods, usually in large neural networks that have more hidden layers, which increases the complexity of the relationship between input and output.

Demand forecasting—Encompasses multiple methods for predicting future demand for a product or service.

Generative adversarial networks (**GANs**)—A method of unsupervised learning using two neural networks in tandem to generate results and then analyze the accuracy of those results.

Hardware—The physical components of a computer system.

Inertial measurement units (IMUs)—Used to measure physical forces, angles, and sometimes magnetic fields around an object. In your phone, the screen rotates based on the phone's physical orientation, which is determined via IMUs that provide information about the phone's position in space.

Machine language—Language that is used to give machines specific instructions about what to do. In contrast to programming languages, humans typically can't read machine languages.

Machine learning—An application of artificial intelligence with the goal of modeling data patterns.

Machine perception—Refers to a machine's ability to take in information from its environment through the use of sensors.

Microelectromechanical systems (MEMS)—Really small electronics made up of components that are 1 to 100 micrometers in size.

Natural language processing (NLP)—Uses artificial intelligence to teach machines to use languages that are spoken and written by humans.

Neural networks—Also referred to as *artificial neural networks*, these are organized in a way that is similar to the way neurons work in the human brain.

Node—In a neural network, a node refers to the computer-based representation of a neuron. *Node* generally refers to a basic unit of a network in computer science. For example, your cell phone is a node in a network of cell phones.

Proximity sensors—Can determine whether something is nearby. Proximity sensors in your phone let your phone know whether your face is next to the screen while you're taking a phone call. This function means your phone doesn't turn off when your face touches it.

Recommender systems—Recommend products and services based on predictions about what a user will like or purchase.

Sensors—Used to give machines the ability to perceive the environment around them.

Software—The part of a computational system that uses machine languages to tell the machine what to do.

Training—The process in which a network or model is learning based on a particular dataset.

Turing test—A well-known test for determining whether a machine has human intelligence.

Unstructured data—Also referred to as *free-form data*, this is data that does not have a set structure (for example, a database) to help machines parse it.