6-2022

An Introduction to Artificial Intelligence

Carlos R. Arias
Seattle Pacific University

2

An Introduction to Artificial Intelligence

Carlos R. Arias

This chapter explores the evolution of artificial intelligence, starting with the first ideas of Alan Turing, going through the promises of its inception, and landing in our current state, when AI invokes a sense of power and awe. Next, the chapter will provide a summary of different technologies related to AI and machine learning, such as deep neural networks, to help the reader distinguish different terminologies. The chapter will end with a discussion of some potential tendencies concerning how AI may be used or evolve in the near future, and some questions about the technology in the long term.

SECTION I: HOW IT CAME TO BE

In the beginning . . .

Since the dawn of human history, we have excelled as creatures because of our distinct abilities. These abilities have grown and evolved through the centuries, so much so that, in the latest decades, we have been able to create machines that can mimic our intelligence, as much as we understand it. Perhaps a new kind of world awaits us.

It is difficult to pinpoint a specific time or event that birthed AI; some might argue that it began with philosophers asking who we are, while others may say that it began with mathematics, or with the advent of computers. For the sake of this summary, we will say it all started with Alan Turing's paper "Computing Machinery and Intelligence," published in 1950, which posed the question of whether machines can perform in such a way that they seem intelligent.[1]

The paper presented what later would be known as the "Turing Test." This test basically places a machine and a human in communication with each other, but the human does not know that they are talking to a machine. If the machine is clever enough to fool the human into believing that the machine is a person, then the machine passes the Turing Test. In other words, the test assesses if "a computer [can] communicate well enough to persuade a human that it, too, is a human."[2]

While the work of Turing and other pioneers may have sparked the concept of AI, it has gone through different seasons over the years. The history of AI can be considered a journey that started with a question: Can a machine think?

Let There Be Light

The inception of the idea of AI, as mentioned, started with the question of if it is possible to create a machine that mimics a human. The Turing Test led to a gathering of scientists at what is known as the "Dartmouth Workshop" in 1956. It was at this conference that the term "artificial intelligence" was coined. This meeting hosted a group of researchers—including John McCarthy, Marvin Minsky, Allen Newel, and Herbert Simon, among others—that would later become known as the frontrunners and founders of AI; McCarthy is considered the father of AI. During this conference, the attendees suggested the idea (some say boldly, others naïvely) 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."[3]

The meeting created a lot of interest and expectations about the field of AI; many embraced the promising possibility of building a thinking machine. During these years of early enthusiasm, research was oriented toward creating software that would be able to solve puzzles and play games; during the 1950s and '60s, game playing was considered a marker of intelligence.

1. Turing, "Computing Machinery and Intelligence," 433–60.
2. Bughin et al., "Artificial Intelligence."
3. Russell et al., "Artificial Intelligence," 17–18.

In addition to games, research was done to create programs that could prove theorems. In the journey of AI, these were the first happy days, funding was flowing, and the field was moving forward.[4]

Walking in the Desert

However, by the late 1960s, the promise of intelligent machines was becoming more an imagined fantasy than an achievable fact. Even though some advancements were made in the early years, they were limited to mimicking how humans perform a task instead of analyzing a problem and then executing a solution. Furthermore, the problems resided in a microworld with limited inputs; as problems got bigger, more real-world-like, they would also become intractable. The field was not producing the practical AI applications that were expected.[5]

The journey stalled. Yet, while it was uphill and on rocky terrain, the trail did not end. Although funding was scarce, research continued. During this time, "expert systems" were born. An expert system is software that encodes the knowledge of one or more experts in a given field. Once the knowledge base is created, the software helps users, normally nonexperts, by providing a course of action based on information it was given about a specific problem. These systems are very focused on specific problems and knowledge bases; to create an expert system, "knowledge engineers" interview experts in a determined area of discourse and map their knowledge into the knowledge base of the system.[6]

In the early years of AI, some of the research concentrated on the simulation of biological neural networks. This research was taken up with energy during the late 1980s. In section III, we will consider the fundamentals of neural networks.

The Rebirth

After the birth of expert systems and the return to neural networks, AI was reborn as a field and started to once again get attention and funding. In addition to advancements in computer science, neuroscience, and other related fields, the improved capabilities of computers had a significant impact

4. Russell et al., "Artificial Intelligence," 18.

5. Russell et al., "Artificial Intelligence," 20; Bughin et al., "Artificial Intelligence."

6. Russell et al., "Artificial Intelligence," 22–23; de Hoog, "Methodologies for Building Knowledge-Based Systems," 8.

in this resurgence. The journey started to find fertile ground and flatter paths to walk. "Machine learning" arose when probability reasoning and other extant technologies were used to create models from experiential data to help to make decisions.

The late 1990s and the first decade of the twenty-first century opened the doors for more actors to join in the journey. The internet became ubiquitous, and the volume of available data started to grow exponentially, leading to "Big Data," referring to the vast amount of data that are produced and accessible. The data come mainly from people; the advent of "Web 2.0" (the ability to produce content instead of just consuming it), social networks, the "Internet of Things," and other technologies created large repositories of data that could be used for machine learning algorithms to create models.

Consider a company that sells products online. Whenever a user is logged in to its website (and perhaps even if they're not), the web server records all the hyperlinks the user clicks during their visit to the site. It can also record how long the user views a given page. Now consider millions of people visiting that company's website. This generates quite a large amount of data that the company can later use to predict the behavior and preferences of its clients. Further, in addition to the information gathered from their website, the company could also mine other websites to gather additional information, such as commentaries on their products or collecting user behavior and preferences.

With the advent of Big Data, the sun started to shine and the fertile ground began to bloom. Computational power increased significantly and there was an abundance of data to train machines. Fields related to AI advanced to provide better tools and understanding, enabling programs to better mimic intelligence. AI was here to stay.

During the 2010s, newer technology developed that depended on the availability of large data centers that provided considerable computing power. "Deep learning" arrived on the scene—the construction of large neural networks enabled the creation of much more advanced models and therefore learn more than the basic neural networks of the past. More details on deep learning will be discussed in section III.

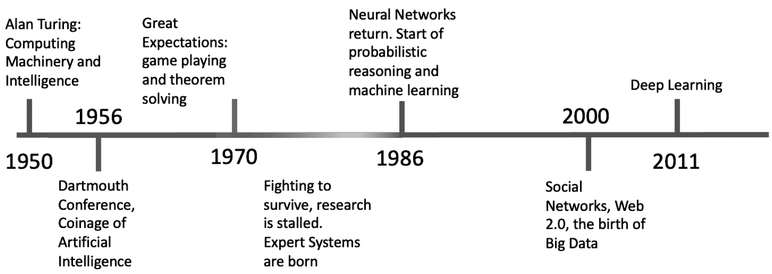

Figure 1: Timeline of the Evolution of AI

Milestones of Artificial Intelligence

The following table presents a very brief summary of some of the milestones of AI. This summary is far from comprehensive but should give the reader an idea of how AI has been evolving and becoming part of our everyday life.

Table 2.1: Milestones of AI	
1956	“Artificial Intelligence” as a term is coined at the Dartmouth Workshop
1966	ELIZA chatbot is created by Joseph Weizenbaum[7]
1986	Self-driving robotic van is created by Ernst Dickmanns[8]
1994	Chinook checkers program wins championship against human[9]
1995	ALICE chatbot developed by Richard Wallace[10]
1997	IBM’s Deep Blue beats world chess champion Garry Kasparov[11]
2000	FDA approves the first robotic surgeon for laparoscopic procedures[12]
2002	Roomba automatic vacuum cleaner is introduced by iRobot Corporation[13]

7. Weizenbaum, “ELIZA—a Computer Program for the Study of Natural Language Communication between Man and Machine,” 36–45.

8. Delcker, “Man Who Invented the Self-Driving Car (in 1986).”

9. Schaeffer et al., “CHINOOK The World Man-Machine Checkers Champion,” 21.

10. Wallace, “Anatomy of A.L.I.C.E.,” 181–210.

11. Goodrich, “How IBM’s Deep Blue Beat World Champion Chess Player Garry Kasparov.”

12. Food And Drug Administration, “FDA Approves New Robotic Surgery Device.”

13. iRobot, “History.”

2011	IBM's Watson computer beats humans in Jeopardy[14]
2011	Apple launches SIRI (Speech Interpretation and Recognition Interface) on iPhone 4[15]
2014	Amazon launches its own speech recognition assistant: Alexa[16]
2017	Google's AlphaGo wins three matches against the world's best Go player Ka Jie[17]
2018	Facebook uses AI to help filter explicit content[18]
2019	AI improves the detection of lung cancer, outperforming radiologists[19]
2020	Waymo driverless taxis begin work in the suburbs of southwest Phoenix[20]

SECTION II: SO, WHAT IS ARTIFICIAL INTELLIGENCE?

"Artificial Intelligence is whatever hasn't been done yet."

—Larry Tesler[21]

Imagine living in the early twentieth century, going about your business as usual, and a person suddenly takes a calculator out of their pocket. Not a fancy scientific calculator. Not a mobile phone. Just a simple add, subtract, multiply, divide calculator. What would you think about it? That it is magical? That it is intelligent?

If someone did the same thing today, most people would look at the device and say that it is too simple, that they can do much more with their mobile phone calculator app. Seems that the calculator lost its genius!

This is one reason why defining AI is so difficult; the concept keeps changing with time, especially as devices get more sophisticated. Smarter even. Further, the concept of intelligence is difficult to define by itself. More on that in later chapters.

14. Best, "IBM Watson."

15. Evans, "WWDC."

16. Lorenzetti, "Forget Siri, Amazon Now Brings You Alexa."

17. Russell, "Google's AlphaGo AI Wins Three-Match Series against the World's Best Go Player."

18. Statt, "Facebook Is Using Billions of Instagram Images to Train Artificial Intelligence Algorithms."

19. Svoboda, "Artificial Intelligence Is Improving the Detection of Lung Cancer," S20–22.

20. Boudway, "Waymo's Self-Driving Future Looks Real Now That the Hype Is Fading."

21. Hofstadter et al., "Eternal Golden Braid," 601.

Let's explore some definitions of AI.

Merriam-Webster:

> 1. A branch of computer science dealing with the simulation of intelligent behavior in computers
> 2. The capability of a machine to imitate intelligent human behavior[22]

Wikipedia:

> In computer science, artificial intelligence (AI), sometimes called machine intelligence, is intelligence demonstrated by machines, unlike the natural intelligence displayed by humans and animals.[23]

Andrew Moore:

> Artificial Intelligence is the science and engineering of making computers behave in ways that, until recently, we thought required human intelligence.[24]

Fred Reed:

> A problem that proponents of AI regularly face is this: When we know how a machine does something "intelligent," it ceases to be regarded as intelligent. If I beat the world's chess champion, I'd be regarded as highly bright.[25]

John McCarthy:

> Artificial Intelligence [is] the ability of certain machines to do things that people are inclined to call intelligent.[26]

The definition of AI depends directly on the definition of intelligence. Different approaches to define intelligence have been taken by researchers in different fields of science.

Russell and Norvig:

> Some have defined intelligence in terms of fidelity to *human* performance, while others prefer an abstract, formal definition

22 *Merriam-Webster*, s.v. "Artificial Intelligence."

23. *Wikipedia*, s.v. "Artificial Intelligence."

24. High, "Carnegie Mellon Dean Of Computer Science On The Future Of AI."

25. "Promise of AI Not so Bright."

26. Bernstein, "Marvin Minsky's Vision of the Future."

> of intelligence called rationality. . . . The subject matter itself also varies: some consider intelligence to be a property of internal *thought processes* and *reasoning*, while others focus on intelligent *behavior*, an external characterization.[27]

In summary it is suggested that depending on the point of view, intelligence is defined somewhere between the dimensions of human versus rational and thought versus behavior. Some think intelligence is an intrinsic ability, while others focus on intelligence as the observable external behaviors of the intelligent entity.

Notice how the basic definitions are based on the concept of "intelligence," but without going into details as to what intelligence actually is; hence the diversity of definitions. Considering the broad nature of the field, having a clear concept of intelligence is challenging. This leads to the range of scientific disciplines that are involved in AI, including computer science, computer engineering, philosophy, psychology, mathematics, neuroscience, linguistics, and biology. See table 2.2.

Table 2.2: Disciplinary Approaches to AI

Discipline	Contributions	Questions
Computer Science	Studies the modeling, creation, and implementation of software that performs the "intelligent" processes.	How can we code this behavior? How can we improve the response time of this program?
Computer Engineering	Studies the construction of the hardware that becomes a computer. It intersects the fields of Computer Science and Electrical Engineering.	How can we build a device that will be sensitive to its environment and be able to respond accordingly? How can we build a device that can be programmed to act "intelligently"?

27. Russell et al., "Artificial Intelligence," 1–2.

Philosophy	Studies knowledge, values, reason, mind, logic, and language, among other issues. It also includes the field of ethics.	How do we reach a logical conclusion given incomplete information? What is knowledge? How is it represented? What is the mind? Is there a relation between intelligence and the mind? What makes humans special? Is a given behavior ethical?
Psychology	Studies the particularities of the mind, reasoning, and behavior.	How can we produce a given behavior? How do people think? What leads people to a certain behavior? How do people learn?
Mathematics (including Statistics)	Studies the numerical relations of objects and abstract representations. Provides numerical tools to represent and operate models.	How can we create a mathematical model to represent knowledge? What can be computed? Is this model robust enough for the given data? What is the probability of success for a given solution of a problem?
Neuroscience	Studies the workings of the brain, including both unconscious and conscious processes.	How does the brain store information? How does the brain process learning?
Linguistics	Studies language and communication, including its symbols, syntax (structure), and semantics (meaning).	How do we communicate? Can we model human communication? Is there a relation between language and thought?

Biology	Studies living things.	How do living things reproduce and evolve? What are the mechanisms of evolution?

Each of these disciplines offer different contributions and questions to the field of AI. Artificial intelligence is inherently multidisciplinary. Many arenas of thought converge in AI, including those above, and more will contribute in the years to come.

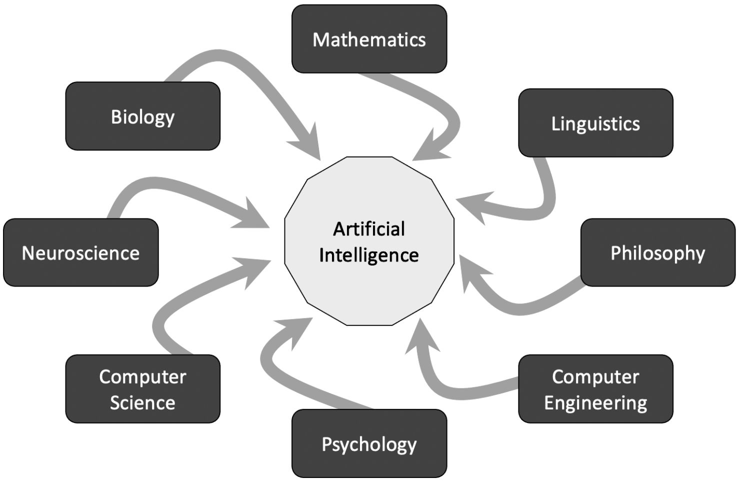

Figure 2: Disciplines That Contribute to AI

Given this multidisciplinary nature, we propose an additional definition of AI:

> Artificial Intelligence is a field of science in which different disciplines (including computer science, computer engineering, philosophy, psychology, mathematics, neuroscience, linguistics, biology, and others) converge with the purpose of creating agents that have the ability to learn (though crudely at the moment), to create models from their learning, to make predictions based on those models and new data, and then to make autonomous decisions (or to help in human decision-making) based on those predictions and any environmental data. All of these characteristics are based on, but not copied from, our current understanding of human intelligence.

Subfields of Artificial Intelligence

Within the field of AI there are many subfields. Keep in mind that AI concerns the creation of intelligence that is not natural, or mimicking the behavior of human intelligence with computer programs. Since intelligent behavior has many aspects, so does AI have many subfields. The major subfields include machine learning, natural language processing, expert systems, computer vision, and robotics. It is important to mention that these subfields are not exclusive; there are often overlaps between them.

Machine learning, as mentioned above, deals with creating programs that consume raw data, build models from that data, and make predictions given new data. To achieve this system of learning, scientists conceive different technologies that somehow resemble our understanding of human learning.

Natural language processing enables computers to recognize human language. Well-known applications of this technology include Cortana, Siri, and Alexa. The main goal of this subfield is to create software and hardware that are able to make sense of human speech and to talk back to users. In the past, it was thought that this was achievable through analysis of the syntax and semantics of human languages. Noah Chomsky, a linguist especially active in the mid- to late twentieth century, even created a grammar to represent human language for programming. However, this sort of natural language modeling turned out to be intractable because human language is intrinsically ambiguous, and natural language processing lost funding. In the late 1980s, a new approach was used involving the application of AI; the increase of computing power and available data has led to machine learning using text analysis to build a language model from user communication.[28]

Expert systems, also mentioned above, represent specific areas of knowledge. These systems are limited to particular disciplines and are built using the knowledge of human experts in that discipline. This knowledge is then programmed into a system to be used by nonexperts. These systems are sometimes thought of as "old-fashioned AI"; however, they do not "learn" by building progressively detailed models, but rather draw from a database "hard-coded" by humans.

Computer vision concerns the processing of images by the computer. It involves mathematical representation of images, machine learning for image recognition, and several other fields like optics and statistics. This subfield is largely known for its advancements in facial and fingerprint

28. Foote, "Brief History of Natural Language Processing (NLP)."

recognition, but it is also used in industry for quality control and in other disciplines including astronomy and history.

Robotics overlaps with other fields of AI. For instance, it uses computer vision to enable robots to "see" using input from visual sensors, machine learning to enable robots to learn from experience (e.g., learning to walk), and natural language processing to enable robots to communicate with humans. Robotics includes the imagining and building of both semi-autonomous and fully autonomous machines (e.g., self-driving cars, self-directing drones), and therefore employs many other forms of technology.

Figure 3 depicts some of the subfields of AI and the relation between them.

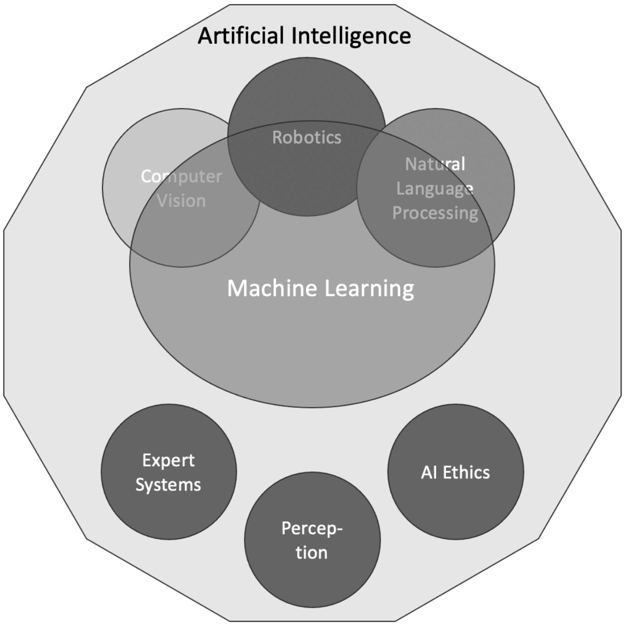

Figure 3: Subfields of AI

SECTION III: THE SCIENCE AND TECHNOLOGY OF ARTIFICIAL INTELLIGENCE

This section provides an overview of the science and technology behind AI. Starting with a general description of machine learning, moving to some types of machine learning, and then looking at an overview of different ways that a machine actually learns. The section ends with a brief discussion of the latest technologies in artificial intelligence and some details about deep neural networks.

Machine Learning

In thinking about AI, a question immediately presents itself: Can a machine learn? Following this question, another quickly emerges: What is learning? It is only by addressing the second question that we can define the extent to which a machine can learn. There are two aspects to the answer to this question in terms of machine learning:

- use available data to create a model to represent what is "learned," and
- set a goal for the computer and then allow it to find the "best" solution to reach that goal based on that model.

Given these two tasks, the process of machine learning (see figure 4) can be outlined as follows.

1. Collect training data that is relevant to what needs to be learned
2. Select an appropriate training algorithm
3. Feed the training algorithm with the training data
4. Retrieve the model created by the training algorithm
5. Use the model to help predict future tendencies or to help in decision making

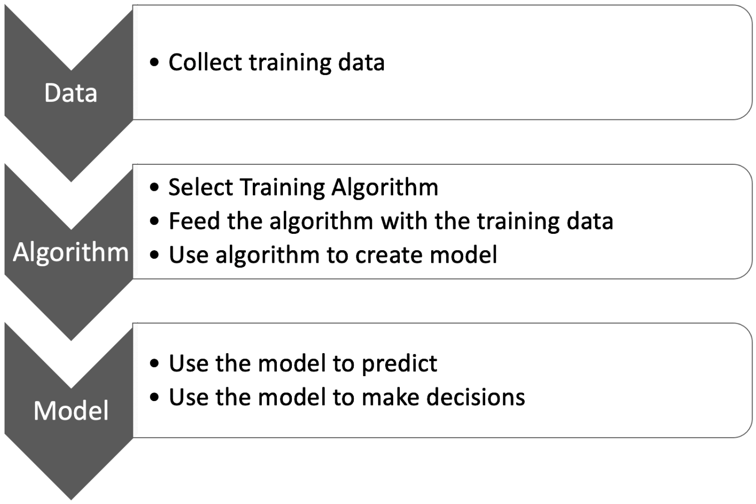

Figure 4: Learning Process Steps

Linear Regression

Statistics is an area of science that has been building representational models long before there was AI. Using inferential statistics, we can create models that predict future outcomes given sufficient historical data. The simplest model of inferential statistics is linear regression.

Linear regression starts with a set of data points in a given space; for simplicity, let's assume the points are in a two-dimensional plane. Figure 5a shows a set of data points that were collected and figure 5b shows how, after applying linear regression to the points, a line appears. This line constitutes the model. Note that, according to figure 5a, we can attest that at time six hours, the traveled distance is around 300 kilometers. However, we are unable to ascertain accurately the traveled distance at 4.5 hours or at hour twelve. Yet the model created through linear regression can help us determine, continuously, all the values of distance given any value of time. Thus, we can predict that at time twelve hours the traveled distance would be almost 700 kilometers.

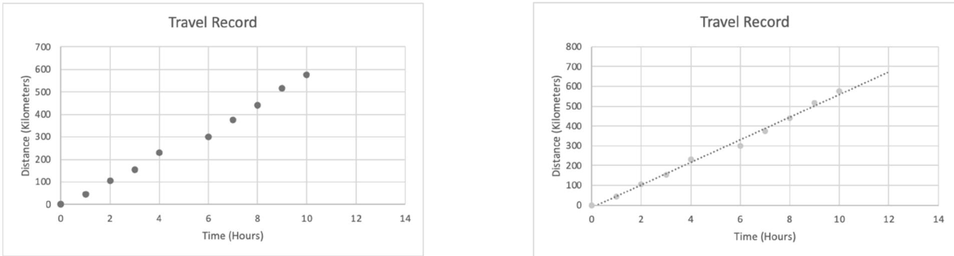

a b

Figure 5: Example of Linear Regression

This simple example highlights a few key concepts.

Concept	Description	Example, from the Travel Record
Training Data	This is a set of data that is used to create the model. As the name implies, the data help train the computer.	In figure 5a, the training data would be the table containing the numbers used to create the graph.
Training Algorithm	The method used to create the appropriate model that best represents the training data.	The algorithm here uses linear regression, producing a line based on the training data that has minimal error point to point.
Model	A mathematical construction that helps predict new values or make decisions.	The generated line shown in figure 5b. The equation that defines that line can predict a distance from a value of time not in the training set.

Clearly this is a very simple model. It assumes that there are only two variables (time and distance), and it assumes a linear behavior between them. Regression can be used for more than two variables, and it can also be used for non-linear behaviors. See figure 6 for an example of non-linear (figure 6a) and three-dimensional data (figure 6b).

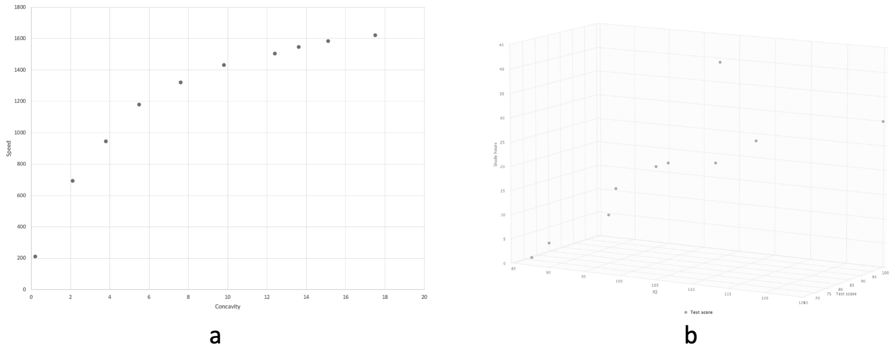

Figure 6: Non-Linear and Multidimensional Data

Classification

Another machine learning technique is classification. The main goal of classification is to create a model that helps classify the features of a subject being studied. For instance, suppose the following data sample.[29]

Day	Outlook	Temperature	Humidity	Wind	Play
1	Sunny	Hot	High	Strong	No
2	Sunny	Hot	High	Weak	No
3	Overcast	Hot	High	Weak	Yes
4	Rain	Mild	High	Weak	Yes
5	Rain	Cool	Normal	Weak	Yes
6	Rain	Cool	Normal	Strong	No
7	Overcast	Cool	Normal	Strong	Yes
8	Sunny	Mild	High	Weak	No
9	Sunny	Cool	Normal	Weak	Yes
10	Rain	Mild	Normal	Weak	Yes
11	Sunny	Mild	Normal	Strong	Yes
12	Overcast	Mild	High	Strong	Yes
13	Overcast	Hot	Normal	Weak	Yes
14	Rain	Mild	High	Strong	No

29. See Witten et al., *Data Mining.*

In the table it can be determined if, depending on the outlook, temperature, humidity, and wind, a person would play or not. Note that this is historic data that reports correlated information. For instance, a person can affirm that, if the outlook is sunny, the temperature mild, the humidity normal, and the wind strong, a person played. A classification algorithm would consume this data and create a model to classify the different conditions of weather to help the user decide whether to play or not.

Some classification algorithms are known as "supervised learning" algorithms. The main characteristic of this type of classification is that the training set contains a correct "answer" for each input.

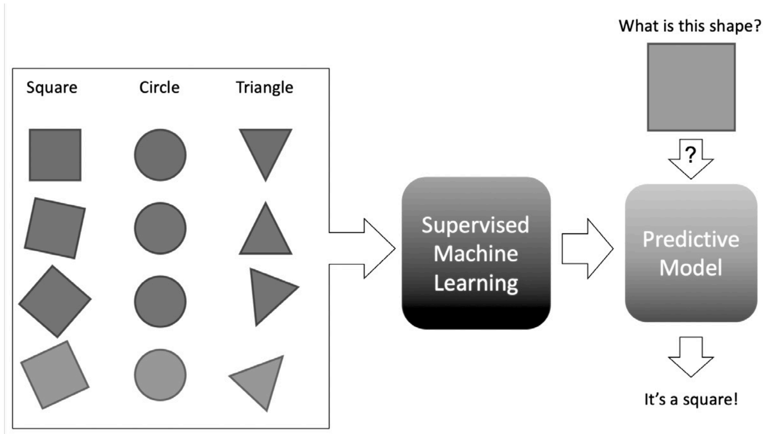

Figure 7: Supervised Learning

Notice that in figure 7, the training set to the left shows twelve shapes. Each shape has a label; for instance, the squares have their label, "square," as do the circles and the triangles. In supervised classification machine learning, the training set has a number of examples with the correct response for each. This training set is fed to a supervised machine learning algorithm to create a predictive model. Once the model is generated, a user can input a shape to the model and the model will classify it according to what it learned from the examples in the training set.

If the training data does not have an "answer" for each of the given examples, an "unsupervised learning" classification algorithm can be used. In figure 8, the shapes are inputted to the machine learning algorithm, but in this case none of the shapes have labels. The algorithm can group these objects according to similarity, but it cannot give them a label that it does not know.

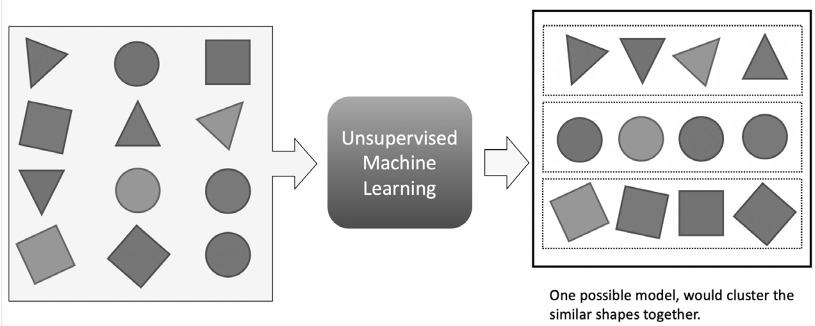

Figure 8: Unsupervised Learning

This type of unsupervised learning is known as "data clustering." In machine learning, the algorithm would create a model by grouping clusters of like data. One application of clustering is "anomaly detection." By clustering the training data, the user can determine outliers of the data and then make decisions based on this discovery.

In addition to learning from training data, sometimes machines can be made to learn from their mistakes. This learning occurs through a feedback loop in which a user evaluates output, which then in turn provides additional input to update the predictive model. This is known as "reinforcement learning" (see figure 9).

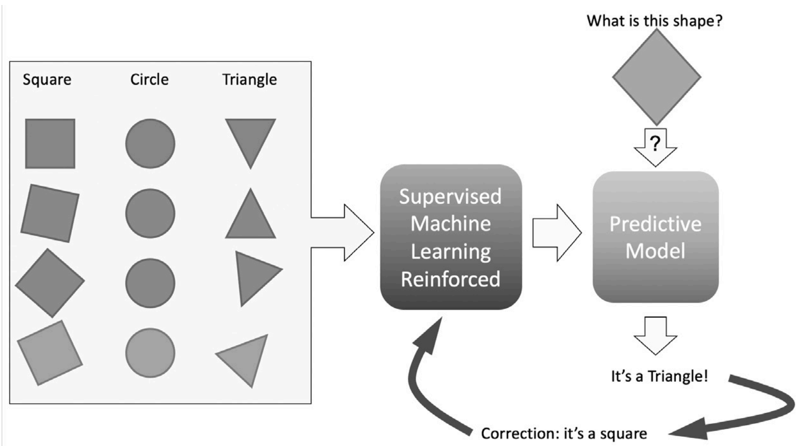

Figure 9: Reinforcement Learning

Association Rules

This machine learning technique creates rules based on historical behavior. For instance, a supermarket can keep a database of all its transactions for the last number of years. It can even associate specific transactions to specific customers—which, incidentally, is why many supermarkets offer customer discount cards. Given this historical data, the computer can be told to analyze it and produce a set of rules: If a customer bought particular items, then it is very likely that the customer will also buy these other particular items. Sometimes these association rules are obvious, but sometimes they are not, in which case they provide unique insight on how consumers behave. Figure 10 illustrates how every transaction, along with the customer information, can be stored in a database. A machine learning algorithm can then use this data to create a model, in this case is a set of association rules that are statistically meaningful according to the recorded data. This technique is also known as "market basket analysis" or "affinity analysis."

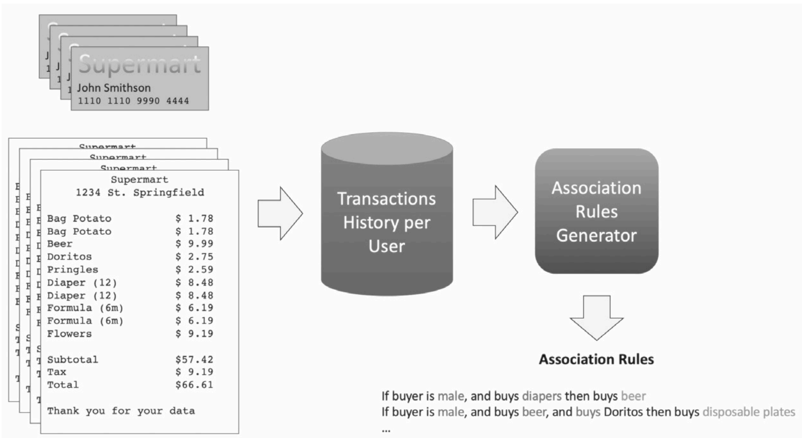

Figure 10: Association Rules Learning

Training Algorithms

So far, our description of machine learning has focused on training algorithms based on linear regression. We now turn to more complex algorithmic technologies used to create representative and predictive models. (See the second step in figure 4.)

Neural Networks

As mentioned above, some of the early AI researchers created a method to train computers based on our current understanding of neurons. The neuron is the basic unit of the nervous system, and it is thought that through networks of neurons humans think and store memories; figure 11 shows a simple schematic of a neuron. The neuron has three distinctive parts: the dendrites, the soma, and the axon. Neurons are a sort of message processor. Messages, in the form of electric impulses, arrive, from several other neurons, into the dendrites. The message is then processed by the soma of the neuron. This processing "decides" whether the neuron should "fire" or not. If it is decided that it should "fire," then a concurrent electrical signal is sent through the axon until it reaches the axon terminals that in turn will stimulate the dendrites of another neuron or neurons.

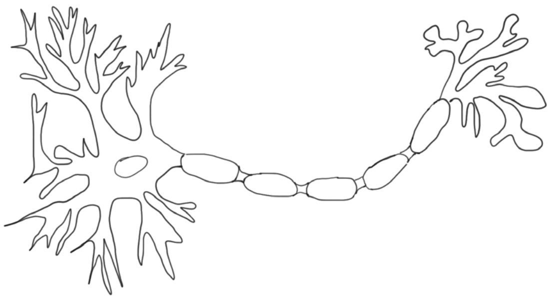

Figure 11: A Neuron

A question comes to mind: When and how does the neuron decide to fire?

Based on the human neural system, computer scientists created a learning model called an artificial neural network.[30] Figure 12a shows a very simplified model of a single neuronal unit, and figure 12b shows a neural network with several interconnected neuronal units. In a neural network, the neuronal unit has several inputs, usually numbers. These values are then multiplied by the "weights" of each artificial dendrite. These weights simulate the importance of that input channel. An example of this process could be the decision of whether or not to go out for dinner based on two considerations: Is it raining? How hungry is the person? Perhaps if it is raining and the person is not that hungry or not hungry at all, then the person will not go out for dinner. However, perhaps if the hunger passes a certain

30. Schmidhuber, "Deep Learning in Neural Networks: An Overview," 85–117.

threshold, then the person will go out for dinner whether it is raining or not. From this example, the inputs could be a report of if it is raining (yes or no) and the level of a person's hunger (not hungry, a little hungry, hungry, or starving), with a higher priority given to hunger and a lower priority given to rain. In an artificial neuronal network, these priorities would be encoded in the weights that multiply the inputs. The resultant values obtained by the prioritization is processed by the neuronal unit, and it decides whether to "fire" or not. If it does, the neuronal unit sends an output that can become the input for another neuronal unit in a network.

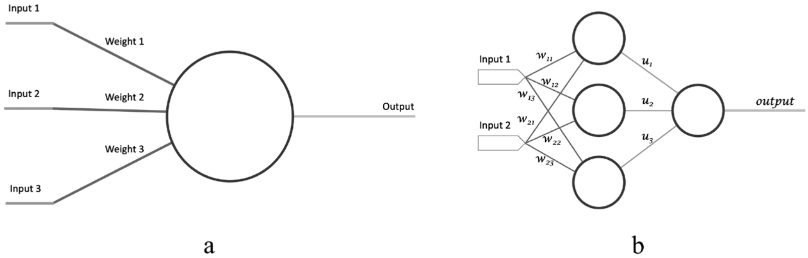

Figure 12: Artificial Neural Networks

Like other machine learning technologies, neural networks need to be trained before they can be used. The training takes in a training set and, based on the data, the algorithm sets the weights for each input pathway for each of the units in the network. This process can be computationally intensive, especially when the network consists in several layers and thousands of units.

Decision Trees

Another model used in machine learning is the decision tree. Like neural networks, this kind of training algorithm consumes a training set and creates a model in the shape of a tree. Each branch poses alternatives and, depending on the chosen alternative, the tree may present more alternatives until a decision is reached. Figure 13 shows a decision tree built with the weather data presented in the classification example above. This decision tree could help a person choose whether or not to play given information about the weather. For instance, if the day is sunny and the humidity is high, the output would suggest not playing that day.

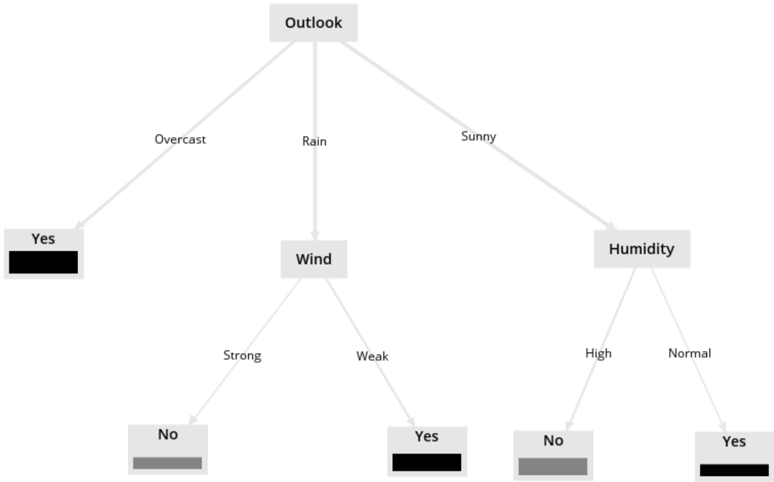

Figure 13: Example of a Decision Tree

Decision trees can be used to make decisions or to find patterns in the data. In medical research, decision trees can be built using clinical information. These trees could then be used by physicians as a diagnostic algorithm[31] or analyzed for patterns that are not obvious. One advantage of decision trees is that they are very intuitive; by looking at them a person can understand how a machine would make a given decision. This is unlike neural networks, in which a person would only see several matrices of numbers representing the weights of connections on each layer of units.

Support Vector Machines

We have seen that one of the simplest methods to create a model is by using linear regression. It was mentioned that the data may not show linear behavior or that it may have more than two dimensions. Sometimes the data does not seem to have a pattern to distinguish the different classes of data. In figure 14a, clearly there is not a line that can separate the gray dots from the black dots. That's where support vector machines come in. This training technique analyzes the data and, using mathematical transformations, finds a non-linear model that represents the data. In figure 14b, a somewhat non-linear separation is found between the different classes of data; the solution

31. Arias et al., "Identification of New Epilepsy Syndromes Using Machine Learning," 1–4.

can be more drastically non-linear as well, as in figure 14c. Support vector machines are particularly powerful in dealing with irregular and multidimensional data.

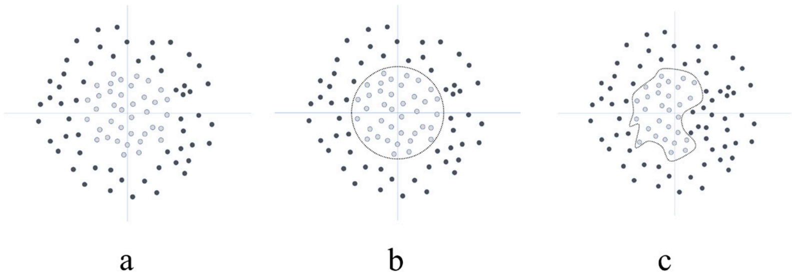

Figure 14: Non-Linear Data

Bayesian Networks

While support vector machines rely on mathematical transformations, Bayesian networks rely on statistical analysis. This machine learning technique uses training data to produce a probabilistic model. Bayesian networks are similar to decision trees, but every decision point can be reached by more than one decision point or input event, and every decision point or input event can lead to more than one conclusion. See the example in figure 15: If it is rush hour, then there is likely going to be a traffic jam; however, there will also be a traffic jam if there is bad weather or when there is a car accident.

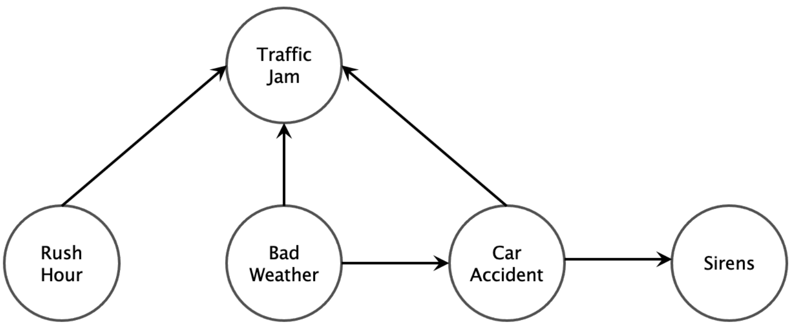

Figure 15: Example of a Bayesian Network

Bayesian networks are built using a statistical tool called "conditional probability." The training algorithm consumes the training data and constructs a model that best represents the data, resulting in a network showing the likelihood of events happening given some other event. Usually, these networks are accompanied with probability tables.

Genetic Algorithms

Natural selection is a process by which living organisms evolve from one state to another. Organisms that are not sufficiently strong or adaptable tend to disappear and therefore do not reproduce. Organisms that have traits that make them stronger and more adaptable survive and are therefore able to pass those characteristics on to the next generation. Natural selection and evolution are concepts that can be also applied to machine learning.

Genetic algorithms are applied to problems that have been intractable. These algorithms start with a first generation of solutions that are evaluated to assess which are "fit" to reproduce. The solutions that make the cut exchange parts (the "crossover" phase) and mutates. This creates a second generation of solutions, and the process repeats, continuing until there is an appropriate solution or the set maximum number of generations has been reached.

Figure 16 shows the process described above. The genetic algorithm starts with a set of solutions (the "initial population") which is then evaluated to find the fittest elements. Some of these elements are discarded, and the remaining ones are crossed over and then mutated. These transformed elements become a new set of possible solutions. If the solution is found, then the process stops. If it is not, then this population is evaluated again, and the process is repeated.

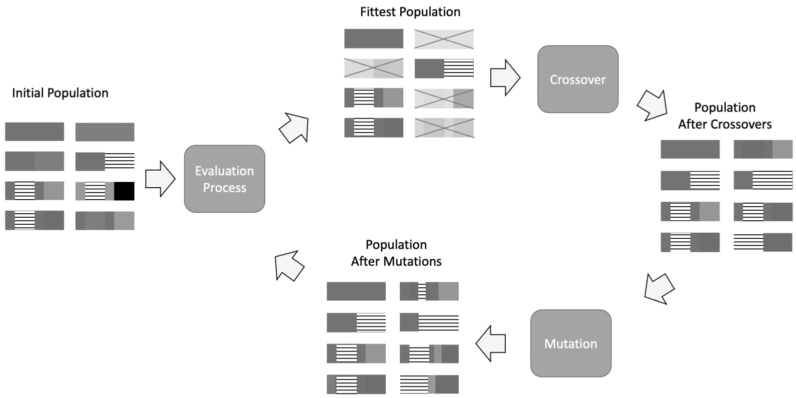

Figure 16: Genetic Algorithm

Clustering Algorithms

Sometimes data do not have labels. This poses a challenge to the classification of the training data. Suppose that there are a thousand images of animals, but none of the images has a label associated with it. This problem calls for an unsupervised machine learning algorithm, one kind of which is called a clustering algorithm. As the name suggests, these algorithms attempt to group data elements for which there is no given label. In figure 17a, all of the data points are black squares, and no other information about them is given except for their position. A clustering algorithm would then try to group these data points together according to some associated characteristic. One possible clustering is shown in figure 17b, with each cluster displaying a different characteristic. In the example above, the clustering algorithm would assess the images of animals and create groups according to the most similar elements. Determining how to define these similarities is one of the challenges of clustering algorithms.

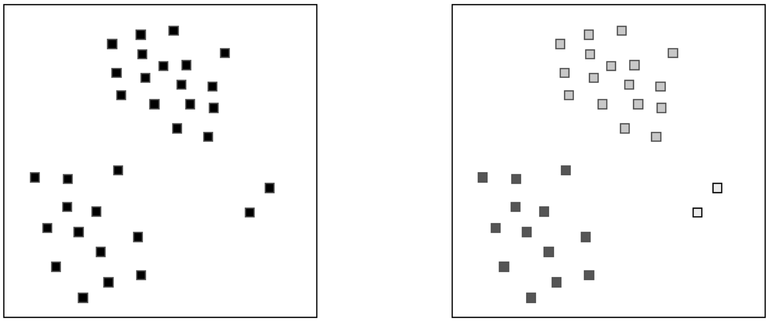

Figure 17: Clustering Data

The topics presented above represent the most basic algorithms for machine learning. Below, we will look at other machine learning and AI algorithms.[32]

Complex Pattern Recognition

Recognizing a face is something we do every day. However, recognizing a moving face, with potential alterations such as glasses, facial hair, different hair cut or color, a face mask, etc., becomes a more interesting problem. Dealing with real-world and real-time data is a significant challenge to pattern recognition algorithms. Examples of complex pattern recognition applications include facial recognition for law enforcement, automatic quality control in production lines, environment analysis for self-driving cars, and medical image analysis.

Information Analysis

One of the skills that every student needs to acquire is the ability to read about a topic and then to provide a summary of it. This certainly represents an activity that many would consider requires intelligence. It is also an activity that AI can accomplish; there are systems that can summarize a given text. Clearly these sorts of systems do not understand the concepts in the text, but the AI is able to compact the text and provide a briefer version of it.

32. Bughin et al., "Artificial Intelligence."

Drawing Conclusions

Artificial intelligence has the ability to draw conclusions based on given knowledge and logical inference. This requires careful representation of the data in such a way that the computer can associate concepts and the creation of a model that maps that representation. Based on that model, the AI then draws conclusions from it. An additional step would involve feeding these conclusions back into the algorithmic model as new data, enabling the AI to learn through "hits and misses."

Forecasting

On any given day, people check forecasts. For instance, many like to know what the weather is going be like, or they want to check the financial markets. By forecasting, we reduce uncertainty intrinsic to future events.

Weather forecasting is a good example of the complexity inherent in learning processes; creating a model to predict the weather has been a challenge given the immense number of variables involved in the atmospheric system. Artificial intelligence has been helpful in this regard. Since the late 1990s, the National Center for Atmospheric Research has been using an AI application called "DICast" to predict the weather.[33] In addition, the National Oceanic and Atmospheric Administration has signed an agreement with Google to "explore the benefits of Artificial Intelligence and Machine Learning" using NOAA's environmental data.[34]

The more complex the system, the more challenging it is to forecast its behavior. Examples of complex systems that AI has attempted to forecast include predicting the electrical energy requirement of a country in real time to balance elements in the power grid[35] and predicting the state of the stock market over different periods of time, even including investor sentiment analysis to increase the precision of the prediction.[36]

33. Haupt et al., "Machine Learning for Applied Weather Prediction," 276–77; "RAL DICast®."

34. Bateman, "AI Agreement to Enhance Environmental Monitoring, Weather Prediction."

35. Motepe et al., "Improving Load Forecasting Process for a Power Distribution Network Using Hybrid AI and Deep Learning Algorithms," 82584–98.

36. Tang et al., "Forecasting Economic Recession through Share Price in the Logistics Industry with Artificial Intelligence (AI)," 70.

Emotion Recognition

Is it possible for computers to feel emotions? This is a very complex philosophical question. At the moment, we would probably say that machines are unable to feel emotions as humans do, but that they are able to recognize some of the features of human emotions. There are three areas in which emotion can be detected: face analysis, text analysis, and voice analysis. In face analysis, the computer is trained to identify different emotions based on the form of the face. Using computer vision, the computer can then determine, with some level of accuracy, the emotion of a person. Scientists are taking this a step further in programming robots to mimic human emotions based on the learning done through facial analysis.[37] Text analysis determines emotions based upon what a person writes; some research suggests it is possible to identify as many as nine emotions from short texts or tweets.[38] This sort of analysis is particularly useful for companies that receive customer feedback, such as in user comments. In addition to facial and text analysis, AI can also determine emotion based on the sound of a person's voice, including in videos.[39]

Deep Learning: Deep Neural Networks

Another term that has become popular in the media in recent years is "deep learning." Deep learning is the result of applying "deep neural networks," which is a neural network that has many layers (usually more than three), each with many neuronal units. For instance, figure 18 shows a neural network that has seven layers of units; some layers have five units and some have seven units. A real-world deep neural network used by Untapt, a company that develops recruitment technology, has sixteen layers with a total of 14 million units.[40] A network like that can take days to train, as it would process several terabytes (thousands of gigabytes) of information.

37. Hall, "How We Feel about Robots That Feel."

38. Argueta et al., "Multilingual Emotion Classifier Using Unsupervised Pattern Extraction from Microblog Data," 1477–502; Saravia et al., "EmoViz," 753–56.

39. Poria et al., "Fusing Audio, Visual, and Textual Clues for Sentiment Analysis from Multimodal Content," 50–59.

40. Untapt, "Talent Science."

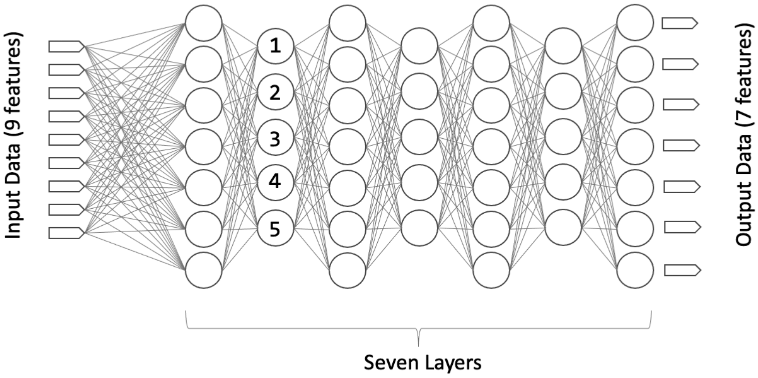

Figure 18: Deep Neural Network

The remarkable ability of deep learning is that it can find patterns in datasets that were thought to be unclassifiable because the problem space was not linear. Further, deep learning can determine features of the patterns it recognizes; each layer of the neural network could be responsible for recognizing a specific aspect of the elements it is analyzing. In this way, the AI "learns" during the training process, gathering its own data rather than relying on information coded by the programmers.

One disadvantage of this technology is that it can be a "black box." It would be very hard for nonspecialists to be able to make sense of the tables of numbers that represent the weights of each of the connections in the neural network. Unlike decision trees that show how they arrive at conclusions, neural networks, and especially deep neural networks, only show conclusions without showing in an intelligible way how they arrived at them.

SECTION IV: WHERE ARTIFICIAL INTELLIGENCE IS WORKING, AND WHAT'S NEXT

Where Can You See Artificial Intelligence Working?

Today, AI is ubiquitous. It can be found throughout industry and everyday activities. It can even be found in the palm of your hand: Siri, Cortana, Google Assistant, Alexa, etc. When you call customer service, many companies will have an automated robot answer. These technologies use natural language processing; AI interprets what you say, decides on the best answer to your question, and responds. It is not perfect, but it is much better than it was even five or ten years ago.

Artificial intelligence can also be found in the stock market. There are many bots (an internet application that performs repetitive tasks) working with stockbrokers to help them decide when to buy and when to sell. These bots learn the behavior of the stock market, make models and projections, and then suggest different alternatives that will increase profits from selected stocks.

In medicine, AI is used to analyze medical images, like MRI or CAT scans. These analyses help physicians with the diagnosis of their patients. Some AI technologies have been shown to detect disease in images that a physician does not.[41]

Artificial intelligence is also used in armament. The US Navy employs an AI technology called the Aegis Combat System, which the Navy describes as "a centralized, automated, command-and-control (C2) and weapons control system that was designed as a total weapon system, from detection to kill."[42] Installed on naval ships, the Aegis automates both defense and attack weaponry. Other AI systems used in the military include enemy detection, navigation, and technologies for missiles or stationary equipment.

Some of the home applications of AI include entertainment and automation. Whatever streaming service a person has, it is likely that the company recommends content based on the movies or shows that the person has watched. These services use machine learning to analyze the viewing history of its users, deducing that if a person watched a given movie, then it is likely that the person will like another particular movie. Media is moving toward even customizing commercials based on user profiles. In terms of home automation, a house could feature AI that controls its environment. For instance, there could be devices to regulate temperature. Unlike an ordinary thermostat that keeps the temperature at a certain level, there are smart thermostats that, to save on heating or cooling costs, adjust the house temperature when they sense that there is no one at home, and they learn when residents usually return home so that they can start getting the house to a comfortable temperature level before the people arrive.

This list of AI applications is not comprehensive; between the writing of these words and your reading of them, there will likely be more advancements making their way into everyday life. What is clear is that there are several participants in AI development, including government, healthcare,

41. Hussain et al., "AI Techniques for COVID-19," 128776–95; Svoboda, "Artificial Intelligence Is Improving the Detection of Lung Cancer," S20–22.

42. Paul Scharre, *Army of None*, 162.

pharmaceutical companies, the military, entertainment industries, global corporations, and academia. This shows how pervasive AI has become.

Latest Advances

- *Deep learning*. As discussed above, deep learning neural networks require several layers with potentially thousands of neuronal units. Training such a network necessitates a considerable amount of electrical energy as this training takes place in data centers with hundreds of computers running sometimes for several days to train a single network. Scientists have found ways to minimize the training process by improving performance, thus saving time and energy. This suggests more applications of deep learning in the near future.
- *Computer code*. Artificial intelligence is now able to write computer code. Given the description of a relatively simple problem, the program can create the code necessary to solve the problem. This technology is still in the early stages, but it is expected that some of the more menial coding tasks will soon be automated.
- *Drug discovery*. Finding medications to treat and cure diseases has always been a challenge. Pharmaceutical companies invest billions of dollars to research drugs in studies that can span several years. Artificial intelligence can help reduce research time by creating models of drugs and simplified models of organisms and then simulating interactions of the two to analyze effects. In addition to these simulations, machine learning is used to gather all available information from previous research and create models to provide suggestions for avenues of future research that were not previously considered.
- *Advanced prosthetics*. Prosthetics have come a long way. For most of human history, any prosthesis consisted merely in the static replacement of a limb. Relatively recently, some movement was incorporated into better prosthetics solutions. Current research is using machine learning to capture the "intention" of user movement by reading nerve impulses. The newest prosthetics are also able to provide biofeedback, giving the user a more realistic experience.
- *COVID-19*. In early 2020, the world began to feel the impacts of the COVID-19 pandemic, a disease caused by the severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2), a virus that first appeared in the Wuhan province in China and then quickly spread throughout the

globe. Scientists around the world began working around the clock to come up with responses, including transmission forecasting, treatment discovery, diagnosis, and prognosis. Artificial intelligence played a role in these solutions thanks to the availability of more clinical data than ever before, including historic clinical information, clinical notes, previous research, data from electronic monitoring and diagnostics equipment, and genetic information. Applying available AI tools, systems were created to model virus transmission and spread. Using drug information and AI, it was possible to determine that some existing drugs could be used to treat COVID-19, saving considerable time in drug research.[43] Artificial intelligence also has aided the significant logistical challenges related to the distribution of the vaccine. In clinical settings, AI is helping with "drug discovery, planning, treatment, and reported outcomes of the COVID-19 patient."[44] It is expected that AI will play an important role in finding further treatments and eventually a cure.[45]

What May Be Coming: A Forecast

- *Healthcare*. It is unlikely that AI will replace physicians. However, advances in AI will continue to impact healthcare profoundly. Physicians will be able to access diagnostic systems that will aid them making decisions. Artificial intelligence will enable the logistics necessary for efficient distribution of healthcare equipment and personnel within an institution or throughout the world. As levels of user familiarity and trust increase, AI will increasingly penetrate the world of healthcare, but it is expected that there will be a human "in the loop" for the foreseeable future.
- *Data generation*. One of the challenges of AI is the amount of data required to train it. Recent cutting-edge research has created techniques for data synthesis using generative adversarial networks (GAN). This technology generates anonymized data that can then be used for training machine learning systems.

43. Hussain et al., "AI Techniques for COVID-19," 128776–95; Pham et al., "Deep Learning Framework for High-Throughput Mechanism-Driven Phenotype Compound Screening and Its Application to COVID-19 Drug Repurposing," 1–11.

44. Mohanty et al., "Application of Artificial Intelligence in COVID-19 Drug Repurposing," 1027–31.

45. Kannan et al., "Role of Artificial Intelligence and Machine Learning Techniques."

- *Artificial intelligence for artificial intelligence*. Can we use AI to create more AI? Researchers have recently used machine learning systems to suggest the training parameters of other machine learning technologies. These systems help with the creation, deployment, management, and operation of AI. Examples include Google AutoML[46] and IBM Auto AI.[47]
- *Manufacturing*. Artificial intelligence continues to integrate into manufacturing industries. It has been decades since automated manufacturing began. In the future, quality control operations, workflow logistics, and management of raw material and finished product will likely become automated. There may come a time when a factory operates with minimal or even no human supervision.
- *Drug discovery*. Artificial intelligence has certainly helped with drug repurposing, which is when drugs that were authorized for use in treatment for a certain disease are discovered to be effective in treatment for a different disease. It is hopeful that AI will also help in the lengthy and expensive process of drug discovery. Advancements in proteomics, genomics, and metabolomics, in conjunction with AI applications may help with new drug creation and determination of drug interactions.

How Far Can It Go?

It is difficult to accurately predict long-term advancement in AI. However, one thing can be said: If a task can be automated, it is very likely that it will be automated. Artificial intelligence will be increasingly present in human life in the years to come, likely in both beneficial and detrimental ways. Developments in healthcare, food and goods manufacturing and distribution, and drug discovery seem promising. However, the creation of automated armaments, AI-driven cyberattacks, and sex robots give more pause. In addition, current applications suffer from inherent biases; the technology is not necessarily inherently biased, but if the data that is used for training is biased, then the models that the AI produces will be biased. This is one more reason to keep humans in the loop as we develop and evaluate our technology.

"Ethical AI," or the ethical consideration of AI, has evolved through several stages. In the first wave, philosophers were consulted to develop

46. "Cloud AutoML Custom Machine Learning Models."
47. IBM, "IBM Watson Studio—AutoAI."

principles and guidelines to mitigate the risks and threats of AI. At the same time, a second wave emerged to develop technical interventions to address fairness and bias in AI, but these efforts did not address the social contexts in which technologies are developed and used. The most recent third wave, described by Carly Kind, is more focused on justice—"social justice, racial justice, economic justice, and environmental justice."[48]

Even though AI can simulate emotions and empathy, it is not generally thought—or desired—that AI will experience feelings as humans do. This implies that there are areas of human activity where AI can help but should not take over, including healthcare, warfare, counseling, and the justice system.

One of the advancements in AI that some are expecting is Artificial General Intelligence (AGI), also known as "The Singularity," which is the name given to the theoretical moment when AI can think for itself and is therefore self-directed. Current AI is only able to solve specific problems in a determinate domain. But Irving J. Good calls AGI the "ultraintelligent machine," with the capacity to solve any problem, including the problem of improving itself, which would lead to an intelligence explosion. Good cautions: "The first ultraintelligent machine is the last invention that man need ever make—provided that the machine is docile enough to tell us how to keep it under control."[49]

Final Remarks

Humans have always had the ability to create tools and techniques. It may have started with the technology of starting and using fire, or tools used for hunting and farming. Most of the time, these technologies have been helpful. Today, we have taken the first steps in creating technologies that can learn. Our success with AI has led to its application in nearly every facet of life, for good and for ill. The tools that we have created, the tools that we create, and the tools that we will create will be tainted by our own imperfections, by our own brokenness. The more power we give to these technologies, the more they will affect us, individually and collectively. That is the reason why, whenever AI directly affects people, there should always be humans overseeing it.

48. Kind, "Term 'Ethical AI' Is Finally Starting to Mean Something."

49. Good, "Speculations concerning the First Ultraintelligent Machine," 31–88.

BIBLIOGRAPHY

Argueta, Carlos, et al. "Multilingual Emotion Classifier Using Unsupervised Pattern Extraction from Microblog Data." *Intelligent Data Analysis* 20 (2016) 1477–502.

Arias, Carlos R., et al. "Identification of New Epilepsy Syndromes Using Machine Learning." *2019 IEEE 39th Central America and Panama Convention*, 2019, 1–4.

"Artificial Intelligence." *Merriam-Webster*. https://www.merriam-webster.com/dictionary/artificial+intelligence.

"Artificial Intelligence." *Wikipedia*. https://en.wikipedia.org/w/index.php?title=Artificial_intelligence&oldid=970134895.

Bateman, John. "AI Agreement to Enhance Environmental Monitoring, Weather Prediction." National Oceanic and Atmospheric Administration, October 27, 2020. https://www.noaa.gov/media-release/ai-agreement-to-enhance-environmental-monitoring-weather-prediction.

Bernstein, Jeremy. "Marvin Minsky's Vision of the Future." *New Yorker*, December 14, 1981. https://www.newyorker.com/magazine/1981/12/14/a-i.

Best, Jo. "IBM Watson: The Inside Story of How the Jeopardy-Winning Supercomputer Was Born, and What It Wants to Do Next." *TechRepublic*, September 9, 2013. https://www.techrepublic.com/article/ibm-watson-the-inside-story-of-how-the-jeopardy-winning-supercomputer-was-born-and-what-it-wants-to-do-next/.

Boudway, Ira. "Waymo's Self-Driving Future Looks Real Now That the Hype Is Fading." *Bloomberg*, January 21, 2021. https://www.bloomberg.com/news/articles/2021–01–21/waymo-self-driving-taxis-are-coming-to-more-u-s-cities.

Bughin, Jacques, et al. "Artificial Intelligence: The Next Digital Frontier?" McKinsey Global Institute, June 2017. https://www.calpers.ca.gov/docs/board-agendas/201801/full/day1/06-technology-background.pdf.

De Hoog, Robert. "Methodologies for Building Knowledge-Based Systems." In *The Handbook of Applied Expert Systems*, edited by Jay Liebowitz, 1–8. New York: CRC, 1998.

Delcker, Janosch. "The Man Who Invented the Self-Driving Car (in 1986)." *Politico*, July 19, 2018. https://www.politico.eu/article/delf-driving-car-born-1986-ernst-dickmanns-mercedes/.

Evans, Jonny. "WWDC: The Evolution of Apple's Siri." *Computerworld*, May 22, 2018. https://www.computerworld.com/article/3275224/wwdc-the-evolution-of-apples-siri.html.

Food and Drug Administration. "FDA Approves New Robotic Surgery Device." *ScienceDaily*, July 17, 2000. https://www.sciencedaily.com/releases/2000/07/000717072719.htm.

Foote, Keith D. "A Brief History of Natural Language Processing (NLP)." DATAVERSITY, May 22, 2019. https://www.dataversity.net/a-brief-history-of-natural-language-processing-nlp/.

Good, Irving John. "Speculations concerning the First Ultraintelligent Machine." *Advances in Computers* 6 (1966) 31–88.

Goodrich, Joanna. "How IBM's Deep Blue Beat World Champion Chess Player Garry Kasparov." *IEEE Spectrum*, January 25, 2021. https://spectrum.ieee.org/the-institute/ieee-history/how-ibms-deep-blue-beat-world-champion-chess-player-garry-kasparov.

Google. "Cloud AutoML." https://cloud.google.com/automl.

Hall, Luisa. "How We Feel about Robots That Feel." *MIT Technology Review*, October 24, 2017. https://www.technologyreview.com/2017/10/24/148259/how-we-feel-about-robots-that-feel/.

Haupt, Sue Ellen, et al. "Machine Learning for Applied Weather Prediction." In 2018 IEEE 14th International Conference on E-Science (e-Science), 276–77. Amsterdam: IEEE, 2018.

High, Peter. "Carnegie Mellon Dean of Computer Science on the Future of AI." *Forbes*, October 30, 2017. https://www.forbes.com/sites/peterhigh/2017/10/30/carnegie-mellon-dean-of-computer-science-on-the-future-of-ai/.

Hofstadter, Douglas R. *Gödel, Escher, Bach: An Eternal Golden Braid*. New York: Basic, 1999.

Hussain, Adedoyin Ahmed, et al. "AI Techniques for COVID-19." *IEEE Access* 8 (2020) 128776–95.

IBM. "IBM Watson Studio—AutoAI." November 6, 2020. https://www.ibm.com/cloud/watson-studio/autoai.

iRobot. "History." https://www.irobot.com/About-iRobot/Company-information/History.

Kannan, Shantani, et al. "The Role of Artificial Intelligence and Machine Learning Techniques: Race for COVID-19 Vaccine." *Archives of Clinical Infectious Diseases* 15 (2020) 10338.

Kind, Carly. "The Term 'Ethical AI' Is Finally Starting to Mean Something." *VentureBeat*, August 23, 2020. https://venturebeat.com/2020/08/23/the-term-ethical-ai-is-finally-starting-to-mean-something/.

Klassner, Frank. "Artificial Intelligence: Introduction." *Crossroads* 3 (1996): 2.

Lorenzetti, Laura. "Forget Siri, Amazon Now Brings You Alexa." *Fortune*, November 11, 2014. https://fortune.com/2014/11/06/forget-siri-amazon-now-brings-you-alexa/.

Mohanty, Sweta, et al. "Application of Artificial Intelligence in COVID-19 Drug Repurposing." *Diabetes and Metabolic Syndrome* 14 (2020) 1027–31.

Motepe, Sibonelo, et al. "Improving Load Forecasting Process for a Power Distribution Network Using Hybrid AI and Deep Learning Algorithms." *IEEE Access* 7 (2019) 82584–98.

Pham, Thai-Hoang, et al. "A Deep Learning Framework for High-Throughput Mechanism-Driven Phenotype Compound Screening and Its Application to COVID-19 Drug Repurposing." *Nature Machine Intelligence* 3 (2021) 247–57.

Poria, Soujanya, et al. "Fusing Audio, Visual, and Textual Clues for Sentiment Analysis from Multimodal Content." *Neurocomputing* 174 (2016) 50–59.

"Promise of AI Not so Bright." *Washington Times*, April 13, 2006. https://www.washingtontimes.com/news/2006/apr/13/20060413–105217–7645r/.

Russell, Jon. "Google's AlphaGo AI Wins Three-Match Series against the World's Best Go Player." *TechCrunch*, May 24, 2017. https://social.techcrunch.com/2017/05/24/alphago-beats-planets-best-human-go-player-ke-jie/.

Russell, Stuart J., and Peter Norvig. *Artificial Intelligence: A Modern Approach*. 4th ed. Pearson Series in Artificial Intelligence. Hoboken, NJ: Pearson, 2021.

Saravia, Elvis, et al. "EmoViz: Mining the World's Interest through Emotion Analysis." *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. New York: ACM, 2015.

Schaeffer, Jonathan, et al. "CHINOOK The World Man-Machine Checkers Champion." *AI Magazine* 17 (1996) 21. https://ojs.aaai.org//index.php/aimagazine/article/view/1208.

Scharre, Paul. *Army of None: Autonomous Weapons and the Future of War*. New York: Norton, 2018.

Schmidhuber, Jürgen. "Deep Learning in Neural Networks: An Overview." *Neural Networks* 61 (2015) 85–117.

Statt, Nick. "Facebook Is Using Billions of Instagram Images to Train Artificial Intelligence Algorithms." *The Verge,* May 2, 2018. https://www.theverge.com/2018/5/2/17311808/facebook-instagram-ai-training-hashtag-images.

Svoboda, Elizabeth. "Artificial Intelligence Is Improving the Detection of Lung Cancer." *Nature* 587 (2020) S20–22.

Tang, Ym, et al. "Forecasting Economic Recession through Share Price in the Logistics Industry with Artificial Intelligence (AI)." *Computation* 8 (2020) 70.

Turing, A. M. "Computing Machinery and Intelligence." *Mind* 59 (1950) 433–60.

UCAR. "RAL DICast®." https://ral.ucar.edu/projects/dicast/.

Untapt, "Talent Science: How Deep Learning in AI Works." https://www.untapt.com/data-science-22222.

Wallace, Richard S. "The Anatomy of A.L.I.C.E." In *Parsing the Turing Test*, edited by Robert Epstein et al., 181–210. Dordrecht, Netherlands: Springer, 2009.

Weizenbaum, Joseph. "ELIZA—a Computer Program for the Study of Natural Language Communication between Man and Machine." *Communications of the ACM* 9 (1966) 36–45.

Witten, I. H., and Eibe Frank. *Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations*. San Francisco: Morgan Kaufmann, 2000.