Introduction to Artificial Intelligence and Machine Learning starts by helping you understand Artificial Intelligence (AI), especially what AI needs to work and why it has failed in the past. You also discover the basis for some of the issues with AI today and how those issues might prove to be nearly impossible to solve in some cases. Of course, along with the issues, you also discover the fixes for some problems and consider where scientists are taking AI in search of answers. For a technology to survive, it must have a group of solid applications that actually work. It also must provide a payback to investors with the foresight to invest in the technology. In the past, AI failed to achieve critical success because it lacked some of these features. AI also suffered from being ahead of its time: true AI needed to wait for the current hardware to actually succeed. Today, you can find AI used in various computer applications and to automate processes.

Mrs. A. Kalaivani, Assistant Professor, Department of Computer Technology, Nallamuthu Gounder Mahalingam College, Pollachi. Doing Ph.D in the field of Data Mining from PSGR Krishnammal College for Women, Coimbatore. Her area of interest is data analysis and predictive models to reveal patterns and trends in data from existing data sources.

9 786205 488263

S. Karpagavalli
A. Kalaivani
R. Jayaprakash

INTRODUCTION TO ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING

A Base of Knowledge to Continue with Your Study of Artificial Intelligence and Machine Learning

LAP LAMBERT Academic Publishing

S. Karpagavalli
A. Kalaivani
R. Jayaprakash

INTRODUCTION TO ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING

S. Karpagavalli
A. Kalaivani
R. Jayaprakash

INTRODUCTION TO ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING

A Base of Knowledge to Continue with Your Study of Artificial Intelligence and Machine Learning

LAP LAMBERT Academic Publishing

Imprint
Any brand names and product names mentioned in this book are subject to trademark, brand or patent protection and are trademarks or registered trademarks of their respective holders. The use of brand names, product names, common names, trade names, product descriptions etc. even without a particular marking in this work is in no way to be construed to mean that such names may be regarded as unrestricted in respect of trademark and brand protection legislation and could thus be used by anyone.

Cover image: www.ingimage.com

Publisher:
LAP LAMBERT Academic Publishing
is a trademark of
Dodo Books Indian Ocean Ltd., member of the OmniScriptum S.R.L Publishing group
str. A.Russo 15, of. 61, Chisinau-2068, Republic of Moldova Europe
Printed at: see last page
ISBN: 978-620-5-48826-3

INTRODUCTION TO
ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING

A Base of Knowledge to Continue with Your Study of Artificial Intelligence and Machine Learning

A. Kalaivani
Dr. S. Karpagavalli
Dr. R. Jayaprakash (Ed)

TABLE OF CONTENTS

LIST OF FIGURES

About This Book

Introduction to Artificial Intelligence and Machine Learningstarts by helping you understand Artificial Intelligence (AI), especially what AI needs to work and why it has failed in the past. You also discover the basis for some of the issues with AI today and how those issues might prove to be nearly impossible to solve in some cases. Of course, along with the issues, you also discover the fixes for some problems and consider where scientists are taking AI in search of answers. For a technology to survive, it must have a group of solid applications that actually work. It also must provide a payback to investors with the foresight to invest in the technology. In the past, AI failed to achieve critical success because it lacked some of these features. AI also suffered from being ahead of its time: true AI needed to wait for the current hardware to actually succeed. Today, you can find AI used in various computer applications and to automate processes. It's also relied on heavily in the medical field and to help improve human interaction. AI is also related to data analysis, machine learning, and deep learning. Sometimes these terms can prove confusing, so one of the reasons to read artificial intelligence is to discover how these technologies interconnect.

AI has a truly bright future today because it has become an essential technology. This book also shows you the paths that AI is likely to follow in the future. The various trends discussed in this book are based on what people are actually trying to do now. The new technology hasn't succeeded yet, but because people are working on it, it does have a good chance of success at some point.

The Purpose of This Book

This book is not meant to be a comprehensive textbook on machine learning. Instead, it will give you a base of knowledge to continue with your study of machine learning and artificial intelligence. In order to continue your studies and master the subject, there is a large degree of studying that must be done. Will discuss the general structure and organization of machine learning models, the common terms, and the basic statistical concepts necessary to use and understand machine learning.

The basic characteristic of machine learning is the use of artificial inductive reasoning. Artificial inductive reasoning means that a specific event gives you cause to generalize a characteristic. This apple is green; therefore, all apples must be green. But here you can see why inductive reasoning on its own is not always perfect, and why it's difficult to train computers to have the same thought process. One given piece of data is not necessarily representative of thousands of other possible pieces of data. Therefore, when we are using statistics and machine learning, we must be using enough data to be able to reason with confidence, without making the wrong inference based on data that is misinterpreted and becomes misleading.

There are things we do every day as humans that we think of as 'common sense.' These types of intuitive decisions cannot be explicitly programmed in a computer, because the variables that help us make our decisions are too difficult to measure. We probably don't need to see a thousand different combinations of chess pieces on a chessboard to think ahead and plan when we are given a situation we haven't seen before. We, as humans, require much fewer data to be able to infer and learn.

1. INTRODUCTION

1.1What s Artificial Intelligence?

The term artificial intelligence stirs emotions. For one thing there is our fascination with intelligence, which seemingly imparts to us humans a special place among life forms. Questions arise such as “What is intelligence?”, “How can one measure intelligence?” or “How does the brain work?”. All these questions are meaningful when trying to understand artificial intelligence. However, the central question for the engineer, especially for the computer scientist, is the question of the intelligent machine that behaves like a person, showing intelligent behavior. The attribute artificial might awaken much different associations. It brings up fears of intelligent cyborgs. It recalls images from science fiction novels. It raises the question of whether our highest good, the soul, is something we should try to understand, model, or even reconstruct.

With such different offhand interpretations, it becomes difficult to define the term artificial intelligence or AI simply and robustly. Nevertheless I would like to try, using examples and historical definitions, to characterize the field of AI. In 1955, John McCarthy, one of the pioneers of AI, was the first to define the term artificial intelligence, roughly as follows:

The goal of AI is to develop machines that behave as though they were intelligent.

To test this definition, the reader might imagine the following scenario. Fifteen or so small robotic vehicles are moving on an enclosed four by four meter square surface. One can observe various behavior patterns. Some vehicles form small groups with relatively little movement. Others move peacefully through the space and gracefully avoid any collision. Still others appear to follow a leader. Aggressive behaviors are also observable. Is what we are seeing intelligent behavior?

According to McCarthy’s definition the aforementioned robots can be described as intelligent. The psychologist Valentin Braitenberg has shown that this seemingly complex behavior can be produced by very simple electrical circuits. So-called Braitenberg vehicles have two wheels, each of which is driven by anindependent electric motor. The speed of each motor is influenced by a light sensoron the front of the vehicle. The more light that hits the sensor,thefaster the motor runs, the vehicle 1 itsconfiguration, moves away from a point light source. Vehicle 2 on the other handmoves toward the light source. Further small modifications can create otherbehavior patterns, such that with these very simple vehicles we can realize theimpressive behavior described above.

Clearly the above definition is insufficient because AI has the goal of solving difficult practical problems which are surely too demanding for the Braitenberg vehicle. In the Encyclopedia Britannica one finds a Definition that goes like:

AI is the ability of digital computers or computer controlled robots to solve problems thatare normally associated with the higher intellectual processing capabilities of humans.

But this definition also has weaknesses. It would admit for example that a computer with large memory that can save a long text and retrieve it on demand displays intelligent capabilities, for memorization of long texts can certainly be considered a higher intellectual processing capability of humans, as can for example the quick multiplication of two 20-digit numbers. According to this definition, then, every computer is an AI system.

Rich, tersely and concisely, characterizes what AI researchers have been doing for the last 50 years. Even in the year 2050, this definition will be up to date. Tasks such as the execution of many computations in a short amount of time are the strong points of digital computers. In this regard they outperform humans by many multiples. In many other areas, however, humans are far superior to machines. For instance, a person entering an unfamiliar room will recognize thesurroundings within fractions of a second and, if necessary, just as swiftly make decisions and plan actions. To date, this task is too demanding for autonomousrobots. According to Rich's definition, this is therefore a task for AI. In fact,research on autonomous robots is an important, current theme in AI. Constructionof chess computers, on the other hand, has lost relevance because they already playat or above the level of grandmasters. It would be dangerous, however, to conclude from Rich's definition that AI is only concerned with the pragmatic implementation of intelligent processes. Intelligent systems, in the sense of Rich's definition, cannot be built without a deep understanding of human reasoning and intelligent action in general, because of which neuroscience is of great importance to AI. This also shows that the other cited definitions reflect important aspects of AI.

A particular strength of human intelligence is adaptivity. We are capable ofadjusting to various environmental conditions and change our behavior accordingly through learning. Precisely because our learning ability is so vastly superior to that of computers, machine learning is, according to Rich's definition, a central subfield of AI.

1.2 Brain Science and Problem Solving

Through research of intelligent systems we can try to understand how the human brain works and then model or simulate it on the computer. Many ideas and principles in the field of neural networks stem from brain science with the related field of neuroscience.

A very different approach results from taking a goal-oriented line of action,starting from a problem and trying to find the most optimal solution. How humans solve the problem is treated as unimportant here. The method, in this approach, is secondary. First and foremost is the optimal intelligent solution to the problem. Rather than employing a fixed method (such as, for example, predicate logic) AI has as its constant goal the creation of intelligent agents for as many different tasks as possible. Because the tasks may be very different, it is unsurprising that the methods currently employed in AI are often also quite different. Similar to medicine,
which encompasses many different, often life-saving diagnostic and therapy procedures, AI also offers a broad palette of effective solutions for widely varying applications. For mental inspiration. Just as in medicine, there is no universal method for all application areas of AI, rather a great number of possible solutions for the great number of various everyday problems,big and small.

Cognitive science is devoted to research into human thinking at a somewhat higher level. Similarly to brain science, this field furnishes practical AI with many important ideas. On the other hand, algorithms and implementations lead to further important conclusions about how human reasoning functions. Thus these three fields benefit from a fruitful interdisciplinary exchange. The subject of this book, however, is primarily problem-oriented AI as a sub-discipline of computer science. There are many interesting philosophical questions surrounding intelligence and artificial intelligence. We humans have consciousness; that is, we can think aboutourselves and even ponder that we are able to think about ourselves. How doesconsciousness come to be? Many philosophers and neurologists now believe thatthe mind and consciousness are linked with matter, that is, with the brain. Thequestion of whether machines could one day have a mind or consciousness couldat some point in the future become relevant. The mind-body problem in particularconcerns whether or not the mind is bound to the body. We will not discussthese questions here.

1.3 The Turing Test and Chatter-bots

Alan Turing made a name for himself as an early pioneer of AI with his definition of an intelligent machine, in which the machine in question must pass the following test. The test person Alice sits in a locked room with two computer terminals. One terminal is connected to a machine, the other with a non-malicious person Bob. Alice can type questions into both terminals. She is given the task of deciding, after five minutes, which terminal belongs to the machine. The machine passes the test if it can trick Alice at least 30% of the time [Tur50]. While the test is very interesting philosophically, for practical AI, which deals with problem solving, it is not a very relevant test. The reasons for this are similar to those mentioned above related to Braitenberg vehicles.

The AI pioneer and social critic Joseph Weizenbaum developed a program named Eliza, which is meant to answer a test subject's questions like a human psychologist. He was in fact able to demonstrate success in many cases. Supposedly his secretary often had long discussions with the program. Today in the internet there are many so-called chatterbots, some of whose initial responses are quite impressive. After a certain amount of time, however, their artificial nature becomes apparent. Some of these programs are actually capable of learning, whileothers possess extraordinary knowledge of various subjects, for example geography or software development. There are already commercial applications for chatterbots in online customer support and there may be others in the field of e-learning. It is conceivable that the learner and the e-learning system could communicate through a chatterbot. The reader may wish to compare several chatterbots and evaluate their intelligence.

1.4 Logic Solves (Almost) All Problems

AI as a practical science of thought mechanization could of course only begin once there were programmable computers. This was the case in the 1950s. Newell and Simon introduced Logic Theorist, the first automatic theorem prover, and thus also showed that with computers, which actually only work with numbers, one can also process symbols. At the same time McCarthy introduced, with the language LISP, a programming language specially created for the processing of symbolic structures. Both of these systems were introduced in 1956 at the historic Dartmouth Conference, which is considered the birthday of AI. In the US, LISP developed into the most important tool for the implementation of symbol-processing AI systems. Thereafter the logical inference rule known as resolution developed into a complete calculus for predicate logic.

In the 1970s the logic programming language PROLOG was introduced as the European counterpart to LISP. PROLOG offers the advantage of allowing direct programming using Horn clauses, a subset of predicate logic. Like LISP, PROLOG has data types for convenient processing of lists.

Until well into the 1980s, a breakthrough spirit dominated AI, especially among many logicians. The reason for this was the string of impressive achievements in symbol processing. With the Fifth Generation Computer Systems project in Japan and the ESPRIT program in Europe, heavy investment went into the construction of intelligent computers.

For small problems, automatic provers and other symbol-processing systems sometimes worked very well. The combinatorial explosion of the search space, however, defined a very narrow window for these successes. This phase of AI was described in [RN10] as the "Look, Ma, no hands!" era. Because the economic success of AI systems fell short of expectations, funding for logic-based AI research in the United States fell dramatically during the 1980s.

1.5 The New Connectionism

During this phase of disillusionment, computer scientists, physicists, and Cognitive scientists were able to show, using computers which were now sufficiently powerful, that mathematically modeled neural networks are capable of learning using training examples, to perform tasks which previously required costly programming. Because of the fault-tolerance of such systems and their ability to recognize patterns, considerable successes became possible, especially in pattern recognition. Facial recognition in photos and handwriting recognition are two exampleapplications. The system Nettalk was able to learn speech from example texts. Under the name connectionism, a new subdiscipline of AI was born.

Connectionism boomed and the subsidies flowed. But soon even here feasibilitylimits became obvious. The neural networks could acquire impressive capabilities,but it was usually not possible to capture the learned concept in simple formulas orlogical rules. Attempts to combine neural nets with logical rules or the knowledgeof human experts met with great difficulties. Additionally, no satisfactory solutionto the structuring and modularization of the networks was found.

1.6 Reasoning under Uncertainty

AI as a practical, goal-driven science searched for a way out of this crisis. One wished to unite logic's ability to explicitly represent knowledge with neural networks' strength in handling uncertainty. Several alternatives were suggested.

The most promising, probabilistic reasoning, works with conditional probabilities for propositional calculus formulas. Since then many diagnostic and expert systems have been built for problems of everyday reasoning using Bayesian networks. The success of Bayesian networks stems from their intuitive comprehensibility, the clean semantics of conditional probability, and from the centuries-old, mathematically grounded probability theory.

The weaknesses of logic, which can only work with two truth values, can be solved by fuzzy logic, which pragmatically introduces infinitely many values between zero and one. Though even today its theoretical foundation is not totally firm, it is being successfully utilized, especially in control engineering.

A much different path led to the successful synthesis of logic and neural networks under the name hybrid systems. For example, neural networks were employed to learn heuristics for reduction of the huge combinatorial search space in proof discovery.

Methods of decision tree learning from data also work with probabilities. Systems like CART, ID3 and C4.5 can quickly and automatically build very accurate decision trees which can represent propositional logic concepts and then be used as expert systems. Today they are a favorite among machine learning techniques.

Since about 1990, data mining has developed as a sub-discipline of AI in the area of statistical data analysis for extraction of knowledge from large databases. Data mining brings no new techniques to AI, rather it introduces the requirement of using large databases to gain explicit knowledge. One application with great market potential is steering ad campaigns of big businesses based on analysis of many millions of purchases by their customers. Typically, machine learning techniques such as decision tree learning come into play here.

1.7 Distributed, Autonomous and Learning Agents

Distributed artificial intelligence, DAI, has been an active area research since about 1985. One of its goals is the use of parallel computers to increase the efficiency of problem solvers. It turned out, however, that because of the high computational complexity of most problems, the use of "intelligent" systems is more beneficial than parallelization itself.

A very different conceptual approach results from the development of autonomous software agents and robots that are meant to cooperate like human teams. As with the aforementioned Braitenberg vehicles, there are many cases in which an individual agent is not capable of solving a problem, even with unlimited resources. Only the cooperation of many agents leads to the intelligent behavior or to the solution of a problem. An ant colony or a termite colony is capable of erecting buildings of very high architectural complexity, despite the fact that no single ant comprehends how the whole thing fits together. This is similar to the situation of provisioning bread for a large city like New York . There is no central planning agency for bread, rather there are hundreds of bakers that know their respective areas of the city and bake theappropriate amount of bread at those locations.

Active skill acquisition by robots is an exciting area of current research. There are robots today, for example, that independently learn to walk or to perform various motor skills related to soccer. Cooperative learning of multiple robots to solve problems together is still in its infancy.

1.8 AI Grows Up

The above systems offered by AI today are not a universal recipe, but a workshop with a manageable number of tools for very different tasks. Most of these tools are well-developed and are available as finished software libraries, often with convenient user interfaces. The selection of the right tool and its sensible use in each individual case is left to the AI developer or knowledge engineer. Like any other artisanship, this requires a solid education, which this book is meant to promote. More than nearly any other science, AI is interdisciplinary, for it draws upon interesting discoveries from such diverse fields as logic, operations research, statistics, control engineering, image processing, linguistics, philosophy, psychology, and neurobiology. On top of that, there is the subject area of the particular application. To successfully develop an AI project is therefore not always so simple, but almost always extremely exciting.

2. CONCEPT LEARNING AND AI GENERAL TO SPECIFIC ORDERING

2.1 The AI Revolution

Around the year 2010 after about 25 years of research on neural networks, scientists could start harvesting the fruits of their research. The very powerful deep learning networks can for example learn to classify images with very high arruracy. Since image classification is of crucial importance for all types of smart robots, this initiated the AI revolution which in turn leads to smart self-driving cars and service robots.

2.2 AI and Society

There have been many scientific books and science fiction novels written on all aspects of this subject. Due to great advances in AI research, we have been on the brink of the age of autonomous robots and the Internet of Things since roughly 2005. Thus we are increasingly confronted with AI in everyday life. The reader, who may soon be working as an AI developer, must also deal with the social impact of this work. As an author of a book on AI techniques, I have the crucial task of examining this topic. I would like to deal with some particularly important aspects of AI which are of great practical relevance for our lives.

2.3 AI and Transportation

In the past 130 years, automotive industry engineers have made great strides. In Germany, one out of every two people owns their own car. These cars are highly reliable. This makes us very mobile and we use this very convenient mobility in work, everyday life and leisure. Moreover, we are dependent on it. Today, we can not get by without a motor vehicle, especially in rural areas with weak public transportation infrastructure, as for instance in Upper Swabia, where the author and his students live.

The next stage of increased convenience in road transportation is now imminent. In a few years, we will be able to buy electric self-driving cars, i.e. robotic cars, which will autonomously bring us to almost any destination. All passengers in the robotic car would be able to read, work or

sleep during the trip. This is possible on public transit already, but passengers in a robotic car would be able to do this at any time and on any route.
Autonomous vehicles that can operate independently could also travel without passengers. This will lead to yet another increase in convenience: robotic taxis. Via a smartphone app, we will be able to order the optimal taxi, in terms of size and equipment, for any conceivable transportation purpose. We will be able to choose whether wewant to travel alone in the taxi or whetherwe are willing to share a ride with other passengers. We will not need our own car anymore. All associated responsibilities and expenses, such as refueling, technical service, cleaning, searching for parking, buying and selling, garage rent, etc. are void, which saves us money and effort.

Besides the immediate gains in comfort and convenience, robotic cars will offer other significant advantages. For example, according to a McKinsey study [GHZ14], we will need far fewer cars and, above all, far fewer parking places in the era of self-driving cars, which will lead to an immense reduction in resource consumption. According to a Lawrence Berkeley National Laboratory study [GS15], electric self-driving cars will cause a 90% reduction in green house emissions per passenger mile due to the vehicles' energy efficiency and the optimized fit between the vehicle and its purpose. Due to their optimal resource utilization, robotic taxis will be much more environmentally friendly than, for example, heavy buses, which often run at low capacity, especially in rural areas. Overall, robot taxis will contribute dramatically to energy savings and thus, among other things, to a significant improvement in CO_2 and climate problems.
Passenger safety will be much higher than it is today. Experts currently estimate future accident rates between zero and ten percent compared to today. Emotional driving ("road rage"), distracted driving and driving under the influence of drugs and alcohol will no longer exist.

Taxi drivers losing their jobs is often cited as a disadvantage of robotic cars. It is almost certain that there will no longer be taxi drivers from about 2030 onwards, but that is not necessarily a problem. As explained in the previous section, our society just needs to deal with the newly gained productivity properly. In addition to the many advantages mentioned above, robotic cars have two critical problems. Firstly, the so-called rebound effect will nullify at least some of the gains in resource, energy and time savings. Shorter driving times as well as more comfortableand cheaper driving will tempt us to drive more. We can only deal with this problem by rethinking our attitude towards consumption and quality of life. Do we have to use the entire time saved for more activities? Here we are all invited to critical reflection.

Another problem we should take seriously is that the robotic cars will need to be networked. In principle, this gives hackers and terrorists the ability to access and manipulate the vehicles' controls through security holes in their network protocols. If a hacker manages to do this once, he could repeat the attack on a grand scale, potentially bringing entire vehicle fleets to a halt, causing accidents, spying on vehicle occupants, or initiating other criminal actions. Here, as in other areas such as home automation and the Internet of Things, IT security experts will be needed to ensure the highest possible security guarantees using tools of the trade such as cryptographic methods. By the way, improved machine learning algorithms will be useful in detecting hacking attacks.

2.4 Service Robotics

In a few years, shortly after self-driving cars, the next bit of consumption bait on the shelves of electronics stores will be service robots. Recently the Google subsidiary

Figure 3.1. : The assistance robot Marvin, deployed in the AsRoBe research project

Boston Dynamics provided an impressive example in its humanoid robot Atlas. Like the new cars, service robots offer a large gain in comfort and convenience which we would probably like to enjoy. One need only imagine such a robot dutifully cleaning and scrubbing after a party from night until morning without a grumble. Or think of the help that an assistance robot like Marvin, shown in Fig., could provide to the elderly8 or to people with disabilities [SPR+16]. In contrast to the robotic cars, however, these benefits come with costlier trade-offs. Completely new markets would be created, more natural resources and more energy would be consumed, and it is not even certain that people's lives would be simplified by the use of service robots in all areas. One of the first applications for robots like Atlas, developed by Boston Dynamics in contract with Google, will probably be military combat.

It is therefore all the more important that, before these robots come to market, we engage in social discourse on this topic. Science fiction films, such as "Ex Machina" (2015) with its female androids, the chilling "I, Robot" (2004) or the humorous "Robot and Frank" (2012), which depicts the pleasant side of a service robot as an old man's helper, can also contribute to such a discussion.

2.5 Agents

Although the term intelligent agents is not new to AI, only in recent years has it gained prominence through [RN10], among others. Agent denotes rather generally a system that processes information and produces an output from an input. These agents may be classified in many different ways.

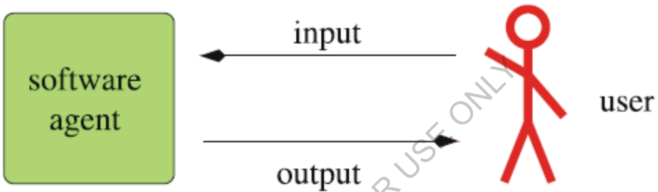

Figure 2.1.: Agent Consists of a Program that Calculates a Result from User Input

In classical computer science, software agents are primarily employed. From (Figure 2.1.) In this case the agent consists of a program that calculates a result from user input.

In robotics, on the other hand, hardware agents (also called autonomous robots) are employed, which additionally have sensors and actuators at their disposal (Figure 2.1.) The agent can perceive its environment with the sensors. With the actuators it carries out actions and changes its environment.

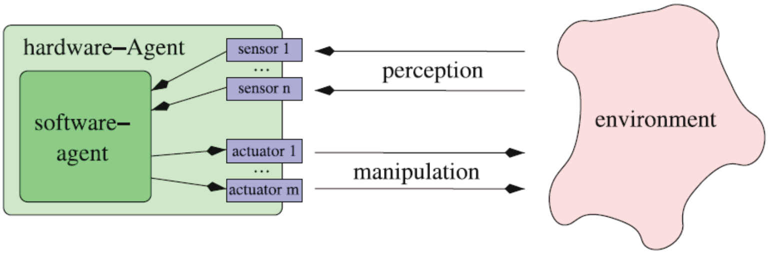

Figure 2.2.: Distinction between Reflex Agents

With respect to the intelligence of the agent, there is a distinction between reflex agents, which only react to input, and agents with memory, which can also include the past in their decisions. For example, a driving robot that through its sensors knows its exact position (and the time) has no way, as a reflex agent, of determining its velocity. If, however, it saves the position, at short, discrete time steps, it can thus easily calculate its average velocity in the previous time interval.

If a reflex agent is controlled by a deterministic program, it represents a function of the set of all inputs to the set of all outputs. An agent with memory, on the other hand, is in general not a function. Reflex agents are sufficient in cases where the problem to be solved involves a Markov Decision process. This is a process in which only the current state is needed to determine the optimal next action

A mobile robot which should move from room 112 to room 179 in a building takes actions different from those of a robot that should move to room 105. In other words, the actions depend on the goal. Such agents are called goal-based.
Example 1.1 A spam filter is an agent that puts incoming emails into wanted or unwanted (spam) categories, and deletes any unwanted emails. Its goal as a goalbased agent is to put all emails in the right category. In the course of this not-so-simple task, the agent can occasionally make mistakes. Because its goal is to classify all emails correctly, it will attempt to make as few errors as possible.

However, that is not always what the user has in mind. Let us compare the following two agents. Out of 1,000 emails, Agent 1 makes only 12 errors. Agent 2 on the other hand makes 38 errors with the same 1,000 emails. Is it therefore worse than Agent 1? The errors of both agents are shown in more detail in the following table, the so-called "confusion matrix":

Agent 1:

		correct class	
		wanted	spam
spam filter decides	wanted	189	1
	spam	11	799

Agent 2:

		correct class	
		wanted	spam
spam filter decides	wanted	200	38
	spam	0	762

Figure 2.3. : Comparison the following Two Agents

Agent 1 in fact makes fewer errors than Agent 2, but those few errors are severe because the user loses 11 potentially important emails. Because there are in this case two types of errors of differing severity, each error should be weighted with the appropriate cost factor.

The sum of all weighted errors gives the total cost caused by erroneous decisions.The goal of a cost-based agent is to minimize the cost of erroneous decisions in the long term, that is, on average.

Analogously, the goal of a utility-based agent is to maximize the utility derived from correct decisions in the long term, that is, on average. The sum of all decisions weighted by their respective utility factors gives the total utility.

Of particular interest in AI are Learning agents, which are capable of changing themselves given training examples or through positive or negative feedback, such that the average utility of their actions grows over time.

The design of an agent is oriented, along with its objective, strongly toward its environment, or alternately its picture of the environment, which strongly depends on it sensors. The environment is observable if the agent always knows the complete state of the world. Otherwise the environment is only partially observable. If an action always leads to the same result, then the environment is deterministic. Otherwise it is nondeterministic. In a discrete environment only finitely many states and actions occur, whereas a continuous environment boasts infinitely manystates or actions.

2.6 Knowledge-Based Systems

An agent is a program that implements a mapping from perceptions to actions. For simple agents this way of looking at the problem is sufficient. For complex applications in which the agent must be able to rely on a large amount of information and is meant to do a difficult task, programming the agent can be very costly and unclear how to proceed. Here AI provides a clear path to follow that will greatly simplify the work.

First we separate knowledge from the system or program, which uses the knowledge to, for example, reach conclusions, answer queries, or come up with a plan. This system is called the inference mechanism. The knowledge is stored in a knowledge base (KB). Acquisition of knowledge in the knowledge base is denoted Knowledge Engineering and is based on various knowledge sources such as human experts, the knowledge engineer, and databases. Active learning systems can also acquire knowledge through active exploration of the world. In Figure. 2.4.

the general architecture of knowledge-based systems is presented.

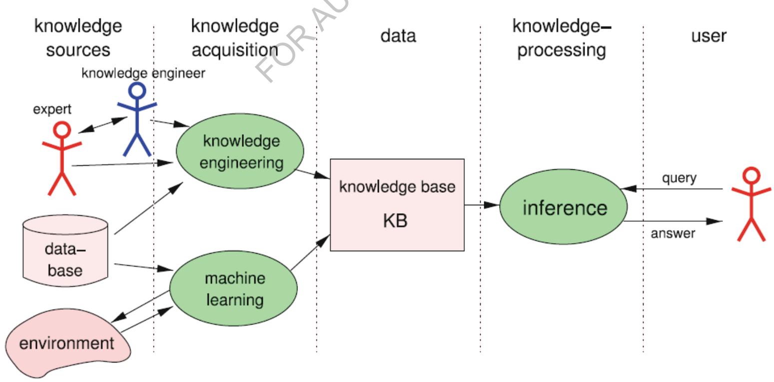

Figure 2.4. Structure of a Classic Knowledge-Processing System

Moving toward a separation of knowledge and inference has several crucial advantages. The separation of knowledge and inference can allow inference systems to be implemented in a largely application-independent way. For example, application of a medical expert system to other diseases is much easier by replacing the knowledge base rather than by programming a whole new system.

Through the decoupling of the knowledge base from inference, knowledge can be stored declaratively. In the knowledge base there is only a description of the knowledge, which is independent from the inference system in use. Without this clear separation, knowledge and processing of inference steps would be interwoven, and any changes to the knowledge would be very costly.

Formal language as a convenient interface between man and machine lends itself to the representation of knowledge in the knowledge base. A powerful language that is accessible by machines and very important in AI.

As an example for a large scale knowledge based system we want to refer to the software agent "Watson". Developed at IBM together with a number of universities, Watson is a question answering program, that can be fed with clues given in natural language. It works on a knowledge base comprising four terabytes of hard disk storage, including the full text of Wikipedia [FNA+09]. Watson was developed within IBM's DeepQA project which is characterized in [Dee11] as follows:

The DeepQA project at IBM shapes a grand challenge in Computer Science that aims to illustrate how the wide and growing accessibility of natural language content and the integration and advancement of Natural Language Processing, Information Retrieval, Machine Learning, Knowledge Representation and Reasoning, and massively parallel computation can drive open-domain automatic Question Answering technology to a point where it clearly and consistently rivals the best human performance.

In the U.S. television quiz show "Jeopardy!", in February 2011, Watson defeated the two human champions Brad Rutter and Ken Jennings in a two-game, combined-point match and won the one million dollar price. One of Watson's particular strengths was its very fast reaction to the questions with the result that Watson often hit the buzzer (using a solenoid) faster than its human competitors and then was able to give the first answer to the question.

The high performance and short reaction times of Watson were due to animplementation on 90 IBM Power 750 servers, each of which contains 32processors, resulting in 2880 parallel processors.

3. NEURAL NETWORKS DEALS WITH DATA SOURCE

3.1 Defining the Role of Data

There is nothing new about data. Every interesting application ever written for a computer has data associated with it. Data comes in many forms — some organized, some not. What has changed is the amount of data. Some people find it almost terrifying that we now have access to so much data that details nearly every aspect of most people's lives, sometimes to a level that even the person doesn't realize. In addition, the use of advanced hardware and improvementsin algorithms make data the universal resource for AI today.

To work with data, you must first obtain it. Today, applications collect data manually, as done in the past, and also automatically, using new methods. However, it's not a matter of just one to two data collection techniques; collection methods take place on a continuum from fully manual to fully automatic.

Raw data doesn't usually work well for analysis purposes. This chapter also helps you understand the need for manipulating and shaping the data so that it meets specific requirements. You also discover the need to define the truth value of the data to ensure that analysis outcomes match the goals set for applications in the first place.

Interestingly, you also have data acquisition limits to deal with. No technology currently exists for grabbing thoughts from someone's mind through telepathic means. Of course, other limits exist, too — most of which you probably already know about but may not have considered.

3.2Finding Data Ubiquitous in This Age

More than a buzzword used by vendors to propose new ways to store data and analyze it, the big data revolution is an everyday reality and a driving force of our times. You may have heard big data mentioned in many specialized scientific and business publications and even wondered what the term really means. From a technical perspective, *big data* refers to large and complex amounts of computer data, so large and intricate that applications can't deal with the data by using additional storage or increasing computer power.

Typical examples of structured data are database tables, in which information is arranged into columns and each column contains a specific type of information. Data is often structured by design. You gather it selectively and record it in its correct place. For example, you might want to place a count of the number of people buying a certain product in a specific column, in a specific table, in a specific database. As with a library, if you know what data you need, you can find it immediately.

Unstructured data consists of images, videos, and sound recordings. You may use an unstructured form for text so that you can tag it with characteristics, such as size, date, or content type. Usually you don't know exactly where data appears in an unstructured dataset because the data appears as sequences of ones and zeros that an application must interpret or visualize.

Transforming unstructured data into a structured form can cost lots of time and effort and can involve the work of many people. Most of the data of the big data revolution is unstructured and stored as it is, unless someone renders it structured.

3.3Using data everywhere

Scientists need more powerful computers than the average person because of their scientific experiments. They began dealing with impressive amounts of data years before anyone coined the term big data. At this point, the Internet didn't produce the vast sums of data that it does today. Remember that big data isn't a fad created by software and hardware vendors but has a basis in many scientific fields, such as astronomy (space missions), satellite (surveillance and monitoring), meteorology, physics (particle accelerators) and genomics (DNA sequences).

Although AI applications can specialize in a scientific field, such as IBM's Watson, which boasts an impressive medical diagnosis capability because it can learn information from millions of scientific papers on diseases and medicine, the actual AI application driver often has more mundane facets. Actual AI applications are mostly prized for being able to recognize objects, move along paths, or understand what people say and to them. Data contribution to the actual AI renaissance that molded it in such a fashion didn't arrive from the classical sources of scientific data.

The Internet now generates and distributes new data in large amounts. Our current daily data production is estimated to amount to about 2.5 quintillion (a number with 18 zeros) bytes, with the lion's share going to unstructured data like videos and audios. All this data is related to common human activities, feelings, experiences, and relations. Roaming through this data, an AI can easily learn how reasoning and acting more human-like works. Here are some examples of the more interesting data you can find:

a)Large repositories of faces and expressions from photos and videos posted on social media websites like Facebook, YouTube, and Google provide information about gender, age, feelings, and possibly sexual preferences, political orientations, or IQ (see https://www.theguardian.com/technology/2017/sep/12/artificial-intelligence-face-recognition-michal-kosinski).

b) Datasets of how people relate to each other and what drives their interest from sources such as social media and search engines. For instance, a study from Cambridge University's Psychometrics Centre claims that Facebook interactions contain a lot of data about intimate relationships (see https:// www.theguardian.com/technology/2015/jan/13/your-computer-knows-youresearchers-cambridge-stanford-university).

Every day, users connect even more devices to the Internet that start storing new personal data. There are now personal assistants that sit in houses, such as Amazon Echo and other integrated smart home devices that offer ways to regulate and facilitate the domestic environment. These are just the tip of the iceberg because many other common tools of everyday life are becoming interconnected (from the refrigerator to the toothbrush) and able to process, record, and transmit data. The Internet of Things (IoT) is becoming a reality. Experts estimate that by 2020, six times as many connected things will exist as there will be people, but research teams and think tanks are already revisiting those figures (http://www.gartner.com/newsroom/id/3165317).

3.4Putting algorithms into action

The human race is now at an incredible intersection of unprecedented volumes of data, generated by increasingly smaller and powerful hardware. The data is also increasingly processed and analyzed by the same computers that the process helped spread and develop. This statement may seem obvious, but data has become so ubiquitous that its value no longer resides only in the information it contains (such as the case of data stored in a firm's database that allows its daily operations), butrather in its use as a means to create new values; such data is described as the "new oil." These new values mostly exist in how applications manicure, store, andretrieve data, and in how you actually use it by means of smart algorithms. Algorithms and AI changed the data game.

AI algorithms have tried different approaches along the way, passing from simple algorithms to symbolic reasoning based on logic and then to expert systems. In recent years, they became neural networks and, in their most mature form, deep learning. As this methodological passage happened, data turned from being the information processed by predetermined algorithms to becoming what molded the algorithm into something useful for the task. Data turned from being just the raw material that fueled the solution to the artisan of the solution itself, as shown in Figure 4,

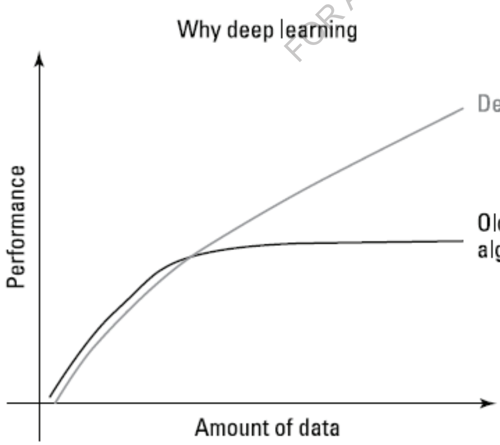

Figure 4.1. : With the present AI solutions, more data equates to more intelligence

Thus, a photo of some of your kittens has become increasingly useful not simply just because of its affective value — depicting your cute little cats — but because it could become part of the learning process of an AI discovering more general concepts, such as what characteristics denote a cat, or understanding what defines cute.

On a larger scale, a company like Google feeds its algorithms from freely available data, such as the content of websites or the text found in publicly available texts and books. Google spider software crawls the web, jumping from website to website, retrieving web pages with their content of text and images. Even if Google gives back part of the data to users as search results, it extracts other kinds of information from the data using its AI algorithms, which learn from it how to achieve other objectives.

3.5 Using Data Successfully

Having plentiful data available isn't enough to create a successful AI. Presently, an AI algorithm can't extract information directly from raw data. Most algorithms rely on external collection and manipulation prior to analysis. When an algorithm collects useful information, it may not represent the right information. The following sections help you understand how to collect, manipulate, and automate data collection from an overview perspective.

3.6 Considering the data sources

The data you use comes from a number of sources. The most common data source is from information entered by humans at some point. Even when a system collects shopping-site data automatically, humans initially enter the information. A human clicks various items, adds them to a shopping cart, specifies characteristics (such as size) and quantity, and then checks out. Later, after the sale, the human gives the shopping experience, product, and delivery method a rating and makes comments. In short, every shopping experience becomes a data collection exercise as well.

Many data sources today rely on input gathered from human sources. Humans also provide manual input. You call or go into an office somewhere to make an appointment with a professional. A receptionist then gathers information from you that's needed for the appointment. This manually collected data eventually ends up in a dataset somewhere for analysis purposes.

Data is also collected from sensors, and these sensors can take almost any form. For example, many organizations base physical data collection, such as the number of people viewing an object in a window, on cellphone detection. Facial recognition software could potentially detect repeat customers.

However, sensors can create datasets from almost anything. The weather service relies on datasets created by sensors that monitor environmental conditions such as rain, temperature, humidity, cloud cover, and so on. Robotic monitoring systems help correct small flaws in robotic operation by constantly analyzing data collected by monitoring sensors. A sensor, combined with a small AI application, could tell you when your dinner is cooked to perfection tonight. The sensor collects data, but the AI application uses rules to help define when the food is properly cooked.

3.7 Obtaining reliable data

The word *reliable* seems so easy to define, yet so hard to implement. Something is reliable when the results it produces are both expected and consistent. A reliable data source produces mundane data that contains no surprises; no one is shocked in the least by the outcome. Depending on your perspective, it could actually be a good thing that most people aren't yawning and then falling asleep when reviewing data. The surprises make the data worth analyzing and reviewing. Consequently, data has an aspect of duality. We want reliable, mundane, fully anticipated data that simply confirms what we already know, but the unexpected is what makes collecting the data useful in the first place.

3.8 Making human input more reliable

Humans make mistakes — its part of being human. In fact, expecting that humans won't make mistakes is unreasonable. Yet, many application designs assume that humans somehow won't make mistakes of any sort. The design expects that everyone will simply follow the rules. Unfortunately, the vast majority of users are guaranteed to not even read the rules because most humans are also lazy or too pressed for time when it comes to doing things that don't reallyhelp them directly.

Consider the entry of a state into a form. If you provide just a text field, some users might input the entire state name, such as Kansas. Of course, some users willmake a typo or capitalization error and come up with Kansus or kANSAS. Settingthese errors, people and organizations have various approaches to performingtasks. Someone in the publishing industry might use the Associated Press (AP)style guide and input Kan. Someone who is older and used to the GovernmentPrinting Office (GPO) guidelines might input Kans. instead. Other abbreviationsare used as well. The U.S. Post Office (USPS) uses KS, but the U.S. Coast Guard usesKA. Meanwhile, the International Standards Organization (ISO) form goes withUS-KS. Mind you, this is just a state entry, which is reasonably straightforward —or so you thought before reading this section. Clearly, because the state isn't goingto change names anytime soon, you could simply provide a drop-down list box onthe form for choosing the state in the required format, thereby eliminating differencesin abbreviation use, typos, and capitalization errors in one fell swoop.

Even with cross-checks and static entries, humans still have plenty of room for making mistakes. For example, entering numbers can be problematic. When a user needs to enter 2.00, you might see 2, or 2.0, or 2., or any of a variety of other entries. Fortunately, parsing the entry and reformatting it will fix the problem, and you can perform this task automatically, without the user's aid.

Unfortunately, reformatting won't correct an errant numeric input. You can partially mitigate such errors by including range checks. A customer can't buy –5 bars of soap. The legitimate way to show the customer returning the bars of soap is to process a return, not a sale. However, the user might have simply made an error, and you can provide a message stating the proper input range for the value.

3.9 Using automated data collection

Some people think that automated data collection solves all the human input issues associated with datasets. In fact, automated data collection does provide a number of benefits:

- Better consistency
- Improved reliability

Lower probability of missing data

Enhanced accuracy

Reduced variance for things like timed inputs

Automated data collection also suffers from both software and hardware errors present in any computing system, but with a higher potential for *soft issues* (which arise when the system is apparently working but isn't providing the desired result) than other kinds of computer-based setups. When the system works, the reliability of the input far exceeds human abilities. However, when soft issues occur, the system often fails to recognize that a problem exists, as a human might, and therefore the dataset could end up containing more mediocre or even bad data

3.10 Manicuring the Data

Some people use the term *manipulation* when speaking about data, giving the impression that the data is somehow changed in an unscrupulous or devious manner. Perhaps a better term would be *manicuring,* which makes the data well shaped and lovely. No matter what term you use, however, raw data seldom meets the requirements for processing and analysis. To get something out of the data, you must manicure it to meet specific needs. The following sections discuss data Manicuring needs.

3.11 Dealing with missing data

To answer a given question correctly, you must have all the facts. You can guess the answer to a question without all the facts, but then the answer is just as likely to be wrong as correct. Often, someone who makes a decision, essentially answering a question, without all the facts is said to jump to a conclusion. When analyzing data, you have probably jumped to more conclusions than you think because of missing data. A *data record,* one entry in a *dataset* (which is all the data), consists of *fields* that contain facts used to answer a question. Each field contains a single kind of data that addresses a single fact. If that field is empty, you don't have the data you need to answer the question using that particular data record.

As part of the process of dealing with missing data, you must know that the data is missing. Identifying that your dataset is missing information can actually be quite hard because it requires you to look at the data at a low level — something that most people aren't prepared to do and is time consuming even if you do have the required skills. Often, your first clue that data is missing is the preposterous answers that your questions get from the algorithm and associated dataset. When the algorithm is the right one to use, the dataset must be at fault.

A problem can occur when the data collection process doesn't include all the data needed to answer a particular question. Sometimes you're better off to actually drop a fact rather than use a considerably damaged fact. If you find that a particular field in a dataset is missing 90 percent or more of its data, the field becomes useless, and you need to drop it from the dataset (or find some way to obtain all that data).

Less damaged fields can have data missing in one of two ways. Randomly missing data is often the result of human or sensor error. It occurs when data records throughout the dataset have missing entries. Sometimes a simple glitch will cause the damage. Sequentially missing data occurs during some type of generalized failure. An entire segment of the data records in the dataset lack the required information, which means that the resulting analysis can become quite skewed.

Fixing randomly missing data is easiest. You can use a simple median or average value as a replacement. No, the dataset isn't completely accurate, but it will likely work well enough to obtain a reasonable answer. In some cases, data scientists used a special algorithm to compute the missing value, which can make the dataset more accurate at the expense of computational time.

Sequentially missing data is significantly harder, if not impossible, to fix because you lack any surrounding data on which to base any sort of guess. If you can find the cause of the missing data, you can sometimes reconstruct it. However, when reconstruction becomes impossible, you can choose to ignore the field. Unfortunately, some answers will require that field, which means that you might need to ignore that particular sequence of data records — potentially causing incorrect output.

3.12 Considering data misalignments

Data might exist for each of the data records in a dataset, but it might not align with other data in other datasets you own. For example, the numeric data in a field in one dataset might be a floating-point type (with decimal point), but an integer type in another dataset. Before you can combine the two datasets, the fields must contain the same type of data.

All sorts of other kinds of misalignment can occur. For example, date fields are notorious for being formatted in various ways. To compare dates, the data formats must be the same. However, dates are also insidious in their propensity for looking the same, but not being the same. For example, dates in one dataset might use Greenwich Mean Time (GMT) as a basis, while the dates in another dataset might use some other time zone. Before you can compare the times, you must align them to the same time zone. It can become even weirder when dates in one dataset come from a location that uses Daylight Saving Time (DST), but dates fromanother location don't.

3.13 Separating useful data from other data

Some organizations are of the opinion that they can never have too much data, but an excess of data becomes as much a problem as not enough. To solve problems efficiently, an AI requires just enough data. Defining the question that you want to answer concisely and clearly helps, as does using the correct algorithm (or algorithm ensemble). Of course, the major problems with having too much data are that finding the solution (after wading through all that extra data) takes longer, and sometimes you get confusing results because you can't see the forest for the trees.

As part of creating the dataset you need for analysis, you make a copy of the original data rather than modify it. Always keep the original, raw data pure so that you can use it for other analysis later. In addition, creating the right data output for analysis can require a number of tries because you may find that the output doesn't meet your needs. The point is to create a dataset that contains only the data needed for analysis, but keeping in mind that the data may need specific kinds of pruning to ensure the desired output.

3.13 Considering the Five Mistruths in Data

Humans are used to seeing data for what it is in many cases: an opinion. In fact, in some cases, people skew data to the point where it becomes useless, a *mistruth.* A computer can't tell the difference between truthful and untruthful data — all it sees is data. One of the issues that make it hard, if not impossible, to create an AI that actually thinks like a human is that humans can work with mistruths and computers can't. The best you can hope to achieve is to see the errant data as outliers and then filter it out, but that technique doesn't necessarily solve the problem because a human would still use the data and attempt to determine a truth based on the mistruths that are there.

Consequently, even machine-derived or sensor-derived data is also subject to generating mistruths that are quite difficult for an AI to detect and overcome. Unfortunately, sensors and other mechanical input methodologies reflect the goals of their human inventors and the limits of what the particular technology is able to detect.

The following sections use a car accident as the main example to illustrate five types of mistruths that can appear in data. The concepts that the accident is trying to portray may not always appear in data and they may appear in different ways than discussed. The fact remains that you normally need to deal with these sorts of things when viewing data.

3.14 Commission

Mistruths of commission are those that reflect an outright attempt to substitute truthful information for untruthful information. For example, when filling out an accident report, someone could state that the sun momentarily blinded them, making it impossible to see someone they hit. In reality, perhaps the person was distracted by something else or wasn't actually thinking about driving (possibly considering a nice dinner). If no one can disprove this theory, the person might get by with a lesser charge. However, the point is that the data would also be contaminated. The effect is that now an insurance company would base premiums on errant data.

Although it would seem as if mistruths of commission are completely avoidable, often they aren't. Human tell "little white lies" to save others embarrassment or to deal with an issue with the least amount of personal effort. Sometimes a mistruth of commission is based on errant input or hearsay. In fact, the sources for errors of commission are so many that it really is hard to come up with a scenario where someone could avoid them entirely. All this said, mistruths of commission are one type of mistruth that someone can avoid more often than not.

3.15 Omission

Mistruths of omission are those where a person tells the truth in every stated fact, but leaves out an important fact that would change the perception of an incident as a whole. Thinking again about the accident report, say that someone strikes a deer, causing significant damage to their car. He truthfully says that the road was wet; it was near twilight so the light wasn't as good as it could be; he was a little late in pressing on the brake; and the deer simply ran out from a thicket at the side of the road. The conclusion would be that the incident is simply an accident.

However, the person has left out an important fact. He was texting at the time. If law enforcement knew about the texting, it would change the reason for the accident to inattentive driving. The driver might be fined and the insurance adjuster would use a different reason when entering the incident into the database. As with the mistruth of commission, the resulting errant data would change how the insurance company adjusts premiums.

Avoiding mistruths of omission is nearly impossible. Yes, someone could purposely leave facts out of a report, but it's just as likely that someone will simply forget to include all the facts. After all, most people are quite rattled after an accident, so it's easy to lose focus and report only those truths that left the most significant impression. Even if a person later remembers additional details and reports them, the database is unlikely to ever contain a full set of truths.

3.16 Perspective

Mistruths of perspective occur when multiple parties view an incident from multiple vantage points. For example, in considering an accident involving a struck pedestrian, the person driving the car, the person getting hit by the car, and a bystander who witnessed the event would all have different perspectives.

An officer taking reports from each person would understandably get different facts from each one, even assuming that each person tells the truth as each knows it. In fact, experience shows that this is almost always the case and what the officer submits as a report is the middle ground of what each of those involved state, augmented by personal experience. In other words, the report will be close to the truth, but not close enough for an AI.

Perspective is perhaps the most dangerous of the mistruths because anyone who tries to derive the truth in this scenario will, at best, end up with an average of the various stories, which will never be fully correct. A human viewing the information can rely on intuition and instinct to potentially obtain a better approximation of the truth, but an AI will always use just the average, which means that the AI is always at a significant disadvantage. Unfortunately, avoiding mistruths of perspective is impossible because no matter how many witnesses you have to the event, the best you can hope to achieve is an approximation of the truth, not the actual truth.

3.17 Bias

Mistruths of bias occur when someone is able to see the truth, but due to personal concerns or beliefs is unable to actually see it. For example, when thinking about an accident, a driver might focus attention so completely on the middle of the road that the deer at the edge of the road becomes invisible. Consequently, the driver has no time to react when the deer suddenly decides to bolt out into the middle of the road in an effort to cross.

A problem with bias is that it can be incredibly hard to categorize. For example, a driver who fails to see the deer can have a genuine accident, meaning that the deer was hidden from view by shrubbery. However, the driver might also be guilty of inattentive driving because of incorrect focus. The driver might also experience a momentary distraction. In short, the fact that the driver didn't see the deer isn't the question; instead, it's a matter of why the driver didn't see the deer. In many cases, confirming the source of bias becomes important when creating an algorithm designed to avoid a biased source.

4. WORKING WITH AI IN HARDWARE APPLICATIONS

4.1 Developing Robots

People often mistake robotics for AI, but robotics are different from AI. Artificial intelligence aims to find solutions to some difficult problems related to human abilities (such as recognizing objects, or understanding speech or text); robotics aims to use machines to perform tasks in the physical world in a partially or completely automated way. It helps to think of AI as the software used to solve problems and of robotics as the hardware for making these solutions a reality.

Robotic hardware may or may not run using AI software. Humans remotely control some robots, as with the da Vinci robot discussed in the "Assisting a surgeon" section of Chapter 7. In many cases, AI does provide augmentation, but the human is still in control. Between these extremes are robots that take abstract orders by humans (such as going from point A to point B on a map or picking up an object) and rely on AI to execute the orders. Other robots autonomously perform assigned tasks without any human intervention. Integrating AI into a robot makes the robot smarter and more useful in performing tasks, but robots don't always need AI to function properly. Human imagination has made the two overlap as a result of sci-fi films and novels.

4.2 Defining Robot Roles

Robots are a relatively recent idea. The word comes from the Czech word *robota,*which means forced labor. The term first appeared in the 1920 play *Rossum'sUniversal Robots,* written by Czech author KarelČapek. However, humanity haslong dreamed of mechanical beings. Ancient Greeks developed a myth of a bronzemechanical man, Talus, built by the god of metallurgy, Hephaestus, at the requestof Zeus, the father of the gods. The Greek myths also contain references toHephaestus building other automata, apart from Talus. *Automata* are selfoperatedmachines that executed specific and predetermined sequences of tasks (as contrasted to robots, which have the flexibility to perform a wide range of tasks). The Greeks actually built water-hydraulic automata that worked the same as an algorithm executed in the physical world. As algorithms, automata incorporate the intelligence of their creator, thus providing the illusion of being self-aware, reasoning machines.

China and Japan have their own versions of automata. Some automata are complex mechanical designs, but others are complete hoaxes, such as the Mechanical Turk, an eighteenth-century machine that was said to be able to play chess but hid a man inside.

4.3 Overcoming the sci-fi view of robots

The first commercialized robot, the Unimate (https://www.robotics.org/joseph-engelberger/unimate.cfm), appeared in 1961. It was simply a robotic arm — a programmable mechanical arm made of metal links and joints — with an end that could grip, spin, or weld manipulated objects according to instructions set by human operators. It was sold to General Motors to use in the production of automobiles. The Unimate had to pick up die-castings from the assembly line and weld them together, a physically dangerous task for human workers. To get an idea of the capabilities of such a machine, check out this video: https://www.youtube.com/watch?v=hxsWeVtb-JQ. The following sections describe the realities of robots today.

4.4 Considering robotic laws

Before the appearance of Unimate, and long before the introduction of many other robot arms employed in industry that started working with human workers in assembling lines, people already knew how robots should look, act, and even think. Isaac Asimov, an American writer renowned for his works in science fiction and popular science, produced a series of novels in the 1950s that suggested a completely different concept of robots from those used in industrial settings.

4.5 Defining actual robot capabilities

Not only are existing robot capabilities still far from the human-like robots found in Asimov's works, they're also of different categories. The kind of biped robot imagined by Asimov is currently the rarest and least advanced.

The most frequent category of robots is the robot arm, such as the previously described Unimate. Robots in this category are also called *manipulators.* You can find them in factories, working as industrial robots, where they assemble and weld at a speed and precision unmatched by human workers. Some manipulators also appear in hospitals to assist in surgical operations. Manipulators have a limited range of motion because they integrate into their location (they might be able to move a little, but not a lot because they lack powerful motors or require an electrical hookup), so they require help from specialized technicians to move to a new location. In addition, manipulators used for production tend to be completely automated (in contrast to surgical devices, which are remote controlled, relying on the surgeon to make medical operation decisions). More than one million manipulators appear throughout the world, half of them located in Japan.

The last kind of robots is the *mobile manipulator,* which can move (as do mobile robots) and manipulate (as do robot arms). The pinnacle of this category doesn't simply consist of a robot that moves and has a mechanical arm but also imitates human shape and behavior. The *humanoid robot* is a biped (has two legs) that has a human torso and communicates with humans through voice and expressions. This kind of robot is what sci-fi dreamed of, but it's not easy to obtain.

4.6 Knowing why it's hard to be a humanoid

Human-like robots are hard to develop, and scientists are still at work on them. Not only does a humanoid robot require enhanced AI capabilities to make them autonomous, it also needs to move as we humans do. The biggest hurdle, though, is getting humans to accept a machine that looks like humans. The following sections look at various aspects of creating a humanoid robot.

Creating a robot that walks

Consider the problem of having a robot walking on two legs (*a bipedal robot*). This is something that humans learn to do adeptly and without conscious thought, but it's very problematic for a robot. Four-legged robots balance easily and they don't consume much energy doing so. Humans, however, do consume energy simply by standing up, as well as by balancing and walking. Humanoid robots, like humans, have to continuously balance themselves, and do it in an effective and economic way. Otherwise, the robot needs a large battery pack, which is heavy and cumbersome, making the problem of balance even more difficult.

A robot with wheels can move easily on roads, but in certain situations, you need a human-shaped robot to meet specific needs. Most of the world's infrastructures are made for a man or woman to navigate. The presence of obstacles, such the passage size, or the presence of doors or stairs, makes using differently shaped robots difficult. For instance, during an emergency, a robot may need to enter a nuclear power station and close a valve. The human shape enables the robot to walk around, descend stairs, and turn the valve wheel.

4.7. Overcoming human reluctance: The uncanny valley

Humans have a problem with humanoid robots that look a little too human. In 1970, a professor at the Tokyo Institute of Technology, Masahiro Mori, studied the impact of robots on Japanese society. He coined the term *Bukimi no TaniGenshō,* which translates to *uncanny valley.* Mori realized that the more realistic robots look, the greater affinity humans feel toward them. This increase in affinity remains true until the robot reaches a certain degree of realism, at which point we start disliking them strongly (even feeling revulsion). The revulsion increases until the robot reaches the level of realism that makes them a copy of a human being. You can find this progression depicted in Figure 5 and described in Mori's original paper at: https://spectrum.ieee.org/automaton/robotics/humanoids/the-uncanny-valley.

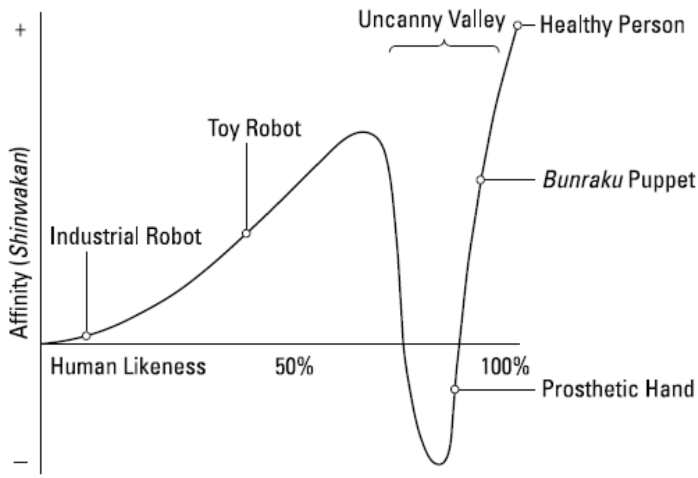

Figure 4.1. : The uncanny valley

Various hypotheses have been formulated about the reasons for the revulsion that humans experience when dealing with a robot that is almost, but not completely, human. Cues that humans use to detect robots are the tone of the robotic voice, the rigidity of movement, and the artificial texture of the robot's skin. The interesting point in the uncanny valley is that if we need humanoid robots because we want them to assist humans, we must also consider their level of realism and key aesthetic details to achieve a positive emotional response that will allow users to accept robot help. Recent observations show that even robots with little human resemblance generate attachment and create bonds with their users. For instance, many U.S. soldiers report feeling a loss when their small tactical robots for explosive detection and handling are destroyed in action.

Working with robots

Different types of robots have different applications. As humans developed and improved the three classes of robots (manipulator, mobile, and humanoid), new fields of application opened to robotics. It's now impossible to enumerate exhaustively all the existing uses for robots, but the following sections touch on some of the most promising and revolutionary uses.

Providing services

Robots provide other care services, both in private and public spaces. The most famous indoor robot is the Roomba vacuum cleaner, a robot that will vacuum the floor of your house by itself (it's a robotic bestseller, having exceeded 3 million units sold), but there are other service robots to consider as well:

Lawn mowing: An incredible variety of lawn-mowing robots exist; you can find some in your local garden shop.

Assistive robots for elder people are far from offering general assistance the way a real nurse does. Robots focus on critical tasks such as remembering medications, helping patients move from a bed to a wheelchair, checking patient physical conditions, raising an alarm when something is wrong, or simply acting as a companion.

4.8 Venturing into dangerous environments

Robots go where people can't, or would be at great risk if they did. Some robots have been sent into space (with the NASA Mars rovers Opportunity and Curiosity being the most notable attempts), and more will support future space exploration. (Chapter 16 discusses robots in space.) Many other robots stay on earth and are employed in underground tasks, such as transporting ore in mines or generating maps of tunnels in caves. Underground robots are even exploring sewer systems, as Luigi (a name inspired from the brother of a famous plumber in videogames) does. Luigi is a sewer-trawling robot developed by MIT's Senseable City Lab to investigate public health in a place where humans can't go unharmed because of high concentrations of chemicals, bacteria, and viruses.

Robots are even employed where humans will definitely die, such as in nuclear disasters like Three Mile Island, Chernobyl, and Fukushima. These robots remove radioactive materials and make the area safer. High-dose radiation even affects robots because radiation causes electronic noise and signal spikes that damage circuits over time. Only *radiation hardened electronic components* allow robots to resist the effects of radiation enough to carry out their job, such as the Little Sunfish, a underwater robot that operates in one of Fukushima's flooded reactors where the meltdown happened.

4.9 Assembling a Basic Robot

An overview of robots isn't complete without discussing how to build one, given the state of the art, and considering how AI can improve its functioning. The following sections discuss robot basics.

Considering the components

A robot's purpose is to act in the world, so it needs *effectors,* which are moving legs or wheels that provide the *locomotion capability*. It also needs arms and pincers to grip, rotate, translate (modify the orientation outside of rotation), and thus provide *manipulating capabilities.* When talking about the capability of the robot to do something, you may also hear the term *actuator* used interchangeably with effectors. An actuator is one of the mechanisms that compose the effectors, allowing a single movement. Thus, a robot leg has different actuators, such as electric motors or hydraulic cylinders that perform movements like orienting the feet or bending the knee

Acting in the world requires determining the composition of the world and understanding where the robot resides in the world. *Sensors* provide input that reports what's happening outside the robot. Devices like cameras, lasers, sonars, and pressure sensors measure the environment and report to the robot what's going on as well as hint at the robot's location. The robot therefore consists mainly of an organized bundle of sensors and effectors. Everything is designed to work together using an architecture, which is exactly what makes up a robot. (Sensors and effectors are actually mechanical and electronic parts that you can use as stand-alone components in different applications.).

The common internal architecture is made of parallel processes gathered into layers that specialize in solving one kind of problem. Parallelism is important. As human beings, we perceive a single flow of consciousness and attention; we don't need to think about basic functions such as breathing, heartbeat, and food digestion because these processes go on by themselves in parallel to conscious thought. Often we can even perform one action, such as walking or driving, while talking or doing something else (although it may prove dangerous in some situations). The same goes for robots. For instance, in the three-layer architecture, a robot has many processes gathered into three layers, each one characterized by a different response time and complexity of answer:

Reactive: Takes immediate data from the sensors, the channels for the robot's perception of the world, and reacts immediately to sudden problems (for instance, turning immediately after a corner because the robot is going to crash on an unknown wall).
Executive: Processes sensor input data, determines where the robot is in the world (an important function called localization), and decides what action to execute given the requirements of the previous layer, the reactive one, and the following one, the deliberative.
Deliberative: Makes plans on how to perform tasks, such as planning how to go from one point to another and deciding what sequence of actions to perform to pick up an object. This layer translates into a series of requirements for the robot that the executive layer carries out.

Sensing the world

Sensors in detail and presents practical applications to help explain self-driving cars. Many kinds of sensors exist, with some focusing on the external world and others on the robot itself. For example, a robotic arm needs to know how much its arm extended or whether it reached its extension limit. Furthermore, some sensors are active (they actively look for information based on a decision of the robot), while others are passive (they receive the information constantly). Each sensor provides an electronic input that the robot can immediately use or process in order to gain a perception.

Perception involves building a local map of real-world objects and determining the location of the robot in a more general map of the known world. Combining data from all sensors, a process called *sensor fusion,* creates a list of basic facts for the robot to use. Machine learning helps in this case by providing vision algorithms using deep learning to recognize objects and segment images. It also puts all the data together into a meaningful representation using unsupervised machine learning algorithms. This is a task called *low-dimensionalembedding,* which means translating complex data from all sensors into a simple flat map or other representation. Determining a robot's location is called *simultaneous localization and mapping (SLAM), and it is just like when you look at a map to understand where you are in a city.*

Controlling a robot

After sensing provides all the needed information, planning provides the robot with the list of the right actions to take to achieve its objectives. Planning is done programmatically or by using a machine learning algorithm, such as Bayesian networks. Developers are experimenting with using reinforcement learning (machine leaning based on trial and error), but a robot is not a toddler (who also relies on trial and error to learn to walk); experimentation may prove time inefficient, frustrating, and costly in the automatic creation of a plan because the robot can be damaged in the process.

Robots have to operate in environments that are partially unknown, changeable, mostly unpredictable, and in a constant flow, meaning that all actions are chained, and the robot has to continuously manage the flow of information and actions in real time. Being able to adjust to this kind of environment can't be fully predicted or programmed, and such an adjustment requires learning capabilities, which AI algorithms provide more and more to robots.

4.10 Flying with Drones

Drones are mobile robots that move in the environment by flying around. Initially connected to warfare, drones have become a powerful innovation for leisure, exploration, commercial delivery, and much more. However, military development still lurks behind developments and causes concern from many AI experts and public figures who foresee them as possibly unstoppable killing machines.

Flying technology is advanced, so drones are more mature than other mobile robots because the key technology to make them work is well understood. The drones' frontier is to incorporate AI. Moving by flying poses some important limits on what drones can achieve, such as the weight they can carry or the actions they can make when arriving at a destination.

Acknowledging the State of the Art

Drones are mobile robots that fly and have existed for a long time, especially for military uses (where the technology originated). The official military name for such flying machines is Unmanned Aircraft System (UAS). More commonly, the public better knows such mobile robots as "drones" because their sound resembles the male bee, but you won't find the term in many official papers because officials prefer names like UAS; or Unmanned Aerial Combat Vehicles (UACV); or Unmanned Aerial Vehicles (UAV); or even RPA (Remotely Piloted Aircraft).

Defining Uses for Drones

Each kind of drone type has current and futuristic applications, and consequently different opportunities to employ AI. The large and small military drones already have their parallel development in terms of technology, and those drones will likely see more use for surveillance, monitoring, and military action in the field. Experts forecast that military uses will likely extend to personal and commercial drones, which generally use different technology from the military ones. (Some overlap exists, such as Duke University's TIKAD, which actually started life in the hobbyist world.)

Apart from rogue uses of small but cheap and easily customizable drones by insurgentsand terrorists groups. Governments are increasingly interested in smaller drones for urban and indoor combat. Indoor places, like corridors or rooms, are where intervention capabilities of aircraft-size Predator and Reaper military drones are limited (unless you need to take down the entire building). The same goes for scout drones, such as Ravens and Pumas, because these drones are made for the operations on the open battlefield, not for indoor warfare.

Commercial drones are far from being immediately employed from shop shelves onto the battlefield, although they offer the right platform for the military to develop various technologies using them. An important reason for the military to use commercial drones is that off-the-shelf products are mostly inexpensive compared to standard weaponry, making them both easily disposable and employable in swarms comprising large number of them. Easy to hack and modify, they require more protection than their already hardened military counterparts do (their communications and controls could be jammed electronically), and they need the integration of some key software and hardware parts before being effectively deployed in any mission.

Navigating in a closed space requires enhanced abilities to avoid collisions, to get directions without needing a GPS (whose signals aren't easily caught while in a building), and to engage a potential enemy. Moreover, drones would need targeting abilities for reconnaissance (spotting ambushes and threats) and for taking out targets by themselves. Such advanced characteristics aren't found in present commercial technology, and they would require an AI solution developed specifically for the purpose. Military researchers are actively developing the required additions to gain military advantage. Recent developments in nimble deep learning networks installed on a standard mobile phone, such as YOLO.

Seeing drones in nonmilitary roles

Currently, commercial drones don't have a lot to offer in the way of advanced functionality found in military models. A commercial drone could possibly take a snapshot of you and your surroundings from an aerial perspective. However, even with commercial drones, a few innovative uses will become quite common in the near future:

1)Delivering goods in a timely fashion, no matter the traffic (being developed by Google X, Amazon, and many startups)
2)Performing monitoring for maintenance and project management
3)Assessing various kinds of damage for insurance
4)Creating field maps and counting herds for farmers
5)Assisting search-and-rescue operations
6)Providing Internet access in remote, unconnected areas (an idea being developed by Facebook)
7)Generating electricity from high-altitude winds

Having goods delivered by a drone is something that hit the public's attention early, thanks to promotion by large companies. One of the earliest and most recognized innovators is Amazon (which promises that a service, Amazon Prime Air, will become operative soon: https://www.amazon.com/Amazon-Prime-Air/b?node=8037720011).

Google promises a similar service with its Project Wing (http://www.businessinsider.com/project-wing-update-future-googledrone-delivery-project-2017-6?IR=T). However, we may still be years away from having a feasible and scalable air delivery system based on drones.
Even though the idea would be to cut intermediaries in the logistic chain in a profitable way, many technical problems and regulatory ambiguities remain to be solved. Behind the media hype showing drones successfully delivering small parcel and other items, such as pizza or burritos, at target locations in an experimental manner. The truth is that drones can't fly far or carry much

weight. The biggest problem is one of regulating the flights of swarms of drones, all of which need to get an item from one point to another. There are obvious issues, such as avoiding obstacles like power lines, buildings, and other drones; facing bad weather; and finding a suitable spot to land near you. The drones would also need to avoid sensitive air space and meet all required regulatory requirements that aircraft meet.

Drones can become your eyes, providing vision in situations that are too costly, dangerous, or difficult to see by yourself. Remotely controlled or semiautonomous (using AI solutions for image detection or processing sensor data), drones can monitor, maintain, surveil, or search and rescue because they can view any infrastructure from above and accompany and support on-demand human operators in their activities. For instance, drones have successfully inspected power lines, pipelines.

Police forces and first-responders around the world have found drones useful for a variety of activities, from search-and-rescue operations to forest fire detection and localization, and from border patrol missions to crowd monitoring. Police are finding newer ways to use drones.

Agriculture is another important area in which drones are revolutionizing work. Not only can they monitor crops, report progress, and spot problems, but they also apply pesticides or fertilizer only where and when needed, as described by MIT Technology Review (https://www.technologyreview.com/s/526491/agricultural-drones/). Drones offer images that are more detailed and less costly than those of an orbital satellite, and they can be employed routinely to *Precision agriculture* uses AI capabilities for movement, localization, vision, anddetection. Precision agriculture could increase agriculture productivity (healthier crops and more food for everyone) while diminishing costs for intervention (no need to spray pesticides everywhere).

Powering up drones using AI

With respect to all drone applications, whether consumer, business, or military related, AI is both a game enabler and a game changer. AI allows many applications to become feasible or better executed because of enhanced autonomy and coordination capabilities. RaffaelloD’Andrea, a

Canadian/Italian/Swiss engineer, professor of dynamic systems and control at ETH Zurich, and drone inventor, demonstrates drone advances in this video: https://www.youtube.com/watch?v=RCXGpEmFbOw. The video shows how drones can become more autonomous by using AI algorithms. *Autonomy* affects how a drone flies, reducing the role of humans issuing drone commands by automatically handling obstacle detection and allowing safe navigation in complicated areas. *Coordination* implies the ability of drones to work together without a central unit to report to and get instructions from, making drones able to exchange information and collaborate in real-time to complete any task.

Taken to its extreme, autonomy may even exclude any human guiding the drone so that the flying machine can determinate the route to take and execute specific tasks by itself. (Humans issue only high-level orders.) When not driven by a pilot, drones rely on GPS to establish an optimal destination path, but that's possible only outdoors, and it's not always precise. Indoor usage increases the need for precision in flight, which requires increased use of other sensor inputs that help the drone understand *proximity surrounds* (the elements of a building, such as a wall protrusion, that could cause it to crash). The cheapest and lightest of these sensors is the camera that most commercial drones have installed as a default device. But having a camera doesn't suffice because it requires proficiency in processing images using computer vision and deep learning techniques.

4.11 Utilizing the AI-Driven Car

A self-driving car (SD car) is an *autonomous vehicle,* which is a vehicle thatcan drive by itself from a starting point to a destination without humanintervention. Autonomy implies not simply having some tasks automated(such as Active Park Assist demonstrated at https://www.youtube.com/watch?v=xW-MhoLImqg), but being able to perform the right steps to achieveobjectives independently. An SD car performs all required tasks on its own, with a human potentially there to observe (and do nothing else). Because SD cars havebeen part of history for more than 100 years.

For a technology to succeed, it must provide a benefit that people see as necessary and not as easily obtained using other methods. That's why SD cars are so exciting. They offer many things of value, other than just driving. This tells you how SD cars will change mobility in significant ways and helps you understand why this is such a compelling technology.

When SD cars become a bit more common and the world comes to accept them as just a part of everyday life, they will continue to affect society. The chapter helps you understand these issues and why they're important. It answers the question of what it will be like to get into an SD car and assume that the car will get you from one place to another without problems.

Finally, SD cars require many sensor types to perform their task. Yes, in some respects you could group these sensors into those that see, hear, and touch, but that would be an oversimplification. The final section of the chapter helps you understand how the various SD car sensors function and what they contribute to the SD car as a whole.

4.12 Getting a Short History

Developing cars that can drive by themselves has long been part of the futuristic vision provided by sci-fi narrative and film since early experiments in the 1920s with radio-operated cars. The problem with these early vehicles is that they weren't practical; someone had to follow behind them to guide them using a radio controller. Consequently, even though the dream of SD cars has been cultivated for so long, the present projects have little to share with the past other than the vision of autonomy.

The military isn't the only one pushing for autonomous vehicles. For a long time, the automotive industry has suffered from overproduction because it can produce more cars than required by market demand. Market demand is down as a result of all sorts of pressures, such as car longevity. In the 1930s, car longevity averaged 6.75 years, but cars today average 10.8 or more years and allow drivers to drive 250,000 or more miles. The decrease in sales has led some makers to exit the industry or fuse together and form larger companies. SD cars are the silver bullet for the industry, offering a way to favorably reshape market demand and convinceconsumers to upgrade. This necessary technology will result in an increase in the production of a large number of new vehicles.

4.13 Understanding the Future of Mobility

SD cars aren't a disruptive invention simply because they'll radically change how people perceive cars, but also because their introduction will have a significant impact on society, economics, and urbanization. At present, no SD cars are on the road yet — only prototypes. (You may think that SD cars are already a commercial reality, but the truth is that they're all prototypes. Look, for example, at the article at https://www.wired.com/story/uber-self-driving-cars-pittsburgh/ and you see phrases such as *pilot projects* used, which you should translate to mean prototypes that aren't ready for prime time.) Many people believe that SD car introduction will require at least another decade, and replacing all the existing car stock with SD cars will take significantly longer. However, even if SD cars are still in the future, you can clearly expect great things from them, as described in the following sections.

4.14 Getting into a Self-Driving Car

Creating a SD car, contrary to what people imagine, doesn't consist of putting a robot into the front seat and letting it drive the car. Humans perform myriad tasks to drive a car that a robot wouldn't know how to perform. To create a human-like intelligence requires many systems connecting to each other and working harmoniously together to define a proper and safe driving environment. Some efforts are under way to obtain an end-to-end solution, rather than rely on separate AI solutions for each need. The problem of developing an SD car requires solving many single problems and having the individual solutions work effectively together. For example, recognizing traffic signs and changing lanes require separate systems.

End-to-end solution is something you often hear when discussing deep learning's role in AI. Given the power of learning from examples, many problems don't require separate solutions, which are essentially a combination of many minor problems, with each one solved by a different AI solution. Deep learning can solve the problem as a whole by solving examples and providing a unique solution that encompasses all the problems that required separate AI solutions in the past. The problem is that deep learning is limited in its capability to actually perform this task today. A single deep learning solution can work for some problems, but others still require that you combine lesser AI solutions if you want to get a reliable, complete solution.

4.15 Putting all the tech together

Under the hood of an SD car are systems working together according to the robotic paradigm of sensing, planning, and acting. Everything starts at the sensing level, with many different sensors telling the car different pieces of information:

1)The GPS tells where the car is in the world (with the help of a map system), which translates into latitude, longitude, and altitude coordinates.
2)The radar, ultrasound, and lidar devices spot objects and provide data about their location and movements in terms of changing coordinates in space.
3)The cameras inform the car about its surroundings by providing image snapshots in digital format.

Many specialized sensors appear in an SD car. The "Overcoming Uncertainty of Perceptions" section, later in this chapter, describes them at length and discloses how the system combines their output. The system must combine and process the sensor data before the perceptions necessary for a car to operate become useful. Combining sensor data therefore defines different perspectives of the world around the car.

Localization is knowing where the car is in the world, a task mainly done by processing the data from the GPS device. GPS is a space-based satellite navigation system originally created for military purposes. When used for civilian purposes, it has some inaccuracy embedded (so that only authorized personal can use it to its full precision). The same inaccuracies also appear in other systems, such as GLONASS (the Russian navigation system), GALILEO (or GNSS, the European system), or the BeiDou (or BDS, the Chinese system). Consequently, no matter what satellite constellation you use, the car can tell that it's on a certain road, but it can miss the lane it's using (or even end up running on a parallel road). In addition to the rough locationprovided by GPS, the system processes the GPS data with lidar sensor data todetermine the exact position based on the details of the surroundings.

The *detection system* determines what is around the car. This system requires many subsystems, with each one carrying out a specific purpose by using a unique mix of sensor data and processing analysis

1)Lane detection is achieved by processing camera images using image data analysis or deep-learning specialized networks for *image segmentation,* in which an image is partitioned into separated areas labeled by type (that is, road, cars, and pedestrians).
2)Traffic signs and traffic lights detection and classification are achieved by processing images from cameras using deep-learning networks that first spot the image area containing the sign or light and then labeling them with the right type (the type of sign or the color of lights).
3) Combined data from radar, lidar, ultrasound, and cameras help locate external objects and track their movements in terms of direction, speed, and acceleration.
4) Lidar data is mainly used for detecting free space on the road (an unobstructed lane or parking space).

Letting AI into the scene

After the sensing phase, which involves helping the SD car determine where it is and what's going on around it, the planning phase begins. AI fully enters the scene at this point. Planning for an SD cars boils down to solving these specific planning tasks:
Route: Determines the path that the car should take. Because you're in the car to go somewhere specific (well, that's not always true, but it's an assumption that holds true most of the time), you want to reach your destination in the fastest and safest way. In some cases, you also must consider cost. Routing algorithms, which are classic algorithms, are there to help.
»Environment prediction: Helps the car to project itself into the future because it takes time to perceive a situation, decide on a maneuver, and complete it. During the time necessary for the maneuver to take place, other cars could decide to change their position or initiate their own maneuvers, too. When driving, you also try to determine what other drivers intend to do to avoid possible collisions. An SD car does the same thing using machine learning prediction to estimate what will happen next and take the future into account.
Behavior planning: Provides the car's core intelligence. It incorporates the practices necessary to stay on the road successfully: lane keeping; lane changing; merging or entering into a road; keeping distance; handling traffic lights, stop signs and yield signs; avoiding obstacles; and much more. All these tasks are performed using AI, such as an expert system that incorporates many drivers' expertise, or a probabilistic model, such as a Bayesian network, or even a simpler machine learning model.

Understanding it is not just AI

After sensing and planning, it's time for the SD car to act. Sensing, planning, and acting are all part of a cycle that repeats until the car reaches its destination and stops after parking. Acting involves the core actions of acceleration, braking, and steering. The instructions are decided during the planning phase, and the car simply executes the actions with controller system aid, such as the Proportional- Integral-Derivative (PID) controller or Model Predictive Control (MPC), which are algorithms that check whether prescribed actions execute correctly and, if not,immediately prescribe suitable countermeasures.

It may sound a bit complicated, but it's just three systems acting, one after the other, from start to end at destination. Each system contains subsystems that solve a single driving problem, as depicted in Figure 4.2.Using the fastest and most reliable algorithms.

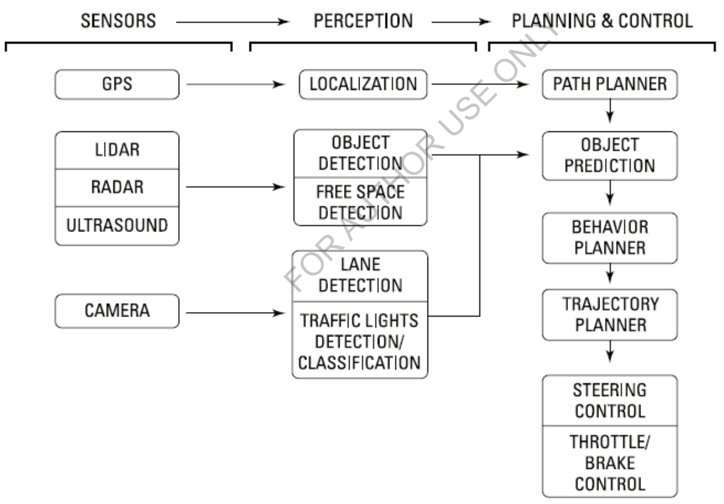

Figure 4.2. : An overall, schematic view of the systems working in an SD car

At the time of writing, this framework is the state of the art. SD cars will likely continue as a bundle of software and hardware systems housing different functions and operations. In some cases, the systems will provide redundant functionality, such as using multiple sensors to track the same external object, or relying on multiple perception processing systems to ensure that you're in the right lane. Redundancy helps ensure zero errors and therefore reduce fatalities. For instance, even when a system like a deep-learning traffic-sign detector fails or is tricked.

4.16 Overcoming Uncertainty of Perceptions

Steven Pinker, professor in the Department of Psychology at Harvard University, says in his book *The Language Instinct: How the Mind Creates Language* that "in robotics, the easy problems are hard and the hard problems are easy." In fact, an AI playing chess against a master of the game is incredibly successful; however, more mundane activities, such as picking up an object from the table, avoiding a collision with a pedestrian, recognizing a face, or properly answering a question over the phone, can prove quite hard for an AI.

The *Moravec paradox* says that what is easy for humans is hard for AI (and vice versa), as explained in the 1980s by robotics and cognitive scientists Hans Moravec, Rodney Brooks, and Marvin Minsk. Humans have had a long time to develop skills such as walking, running, picking up an object, talking, and seeing; these skills developed through evolution and natural selection over millions of years. To survive in this world, humans do what all the living beings have done since life has existed on earth. Conversely, high abstraction and mathematics are a relatively new discovery for humans, and we aren't naturally adapted for them.

Cars have some advantages over robots, which have to make their way in buildings and on outside terrain. Cars operate on roads specifically created for them, usually well-mapped ones, and cars already have working mechanical solutions for moving on road surfaces.

Actuators aren't the greatest problem for SD cars. Planning and sensing are what pose serious hurdles. Planning is at a higher level (what AI generally excels in). When it comes to general planning, SD cars can already rely on GPS navigators, a type of AI specialized in providing directions. Sensing is the real bottleneck for SD cars because without it, no planning and actuation are possible. Drivers sense the road all the time to keep the car in its lane, to watch out for obstacles, and to respect the required rules.

Sensing hardware is updated continuously at this stage of the evolution of SD cars to find more reliable, accurate, and less costly solutions. On the other hand, both processing sensor data and using it effectively rely on robust algorithms, such as the *Kalman filter.*

4.17 Introducing the car's senses

Sensors are the key components for perceiving the environment, and an SD carcan sense in two directions, internal and external:

Proprioceptive sensors: Responsible for sensing vehicle state, such as systems status (engine, transmission, braking, and steering), and the vehicle's position in the world by using GPS localization, rotation of the wheels, the speed of the vehicle, and its acceleration.

Exteroceptive sensors: Responsible for sensing the surrounding environment by using sensors such as camera, lidar, radar, and ultrasonic sensor.

Both proprioceptive and exteroceptive sensors contribute to SD car autonomy. GPS localization, in particular, provides a guess (possibly viewed as a rough estimate) as to the SD car's location, which is useful at a high level for planning directions and actions aimed at getting the SD car to its destination successfully. The GPS helps an SD car in the way it helps any human driver: providing the right directions.

The exteroceptive sensors (shown in Figure 4.3.) help the car specifically in driving. They replace or enhance human senses in a given situation. Each of them offers a different perspective of the environment; each suffers specific limitations; and each excels at different capabilities.

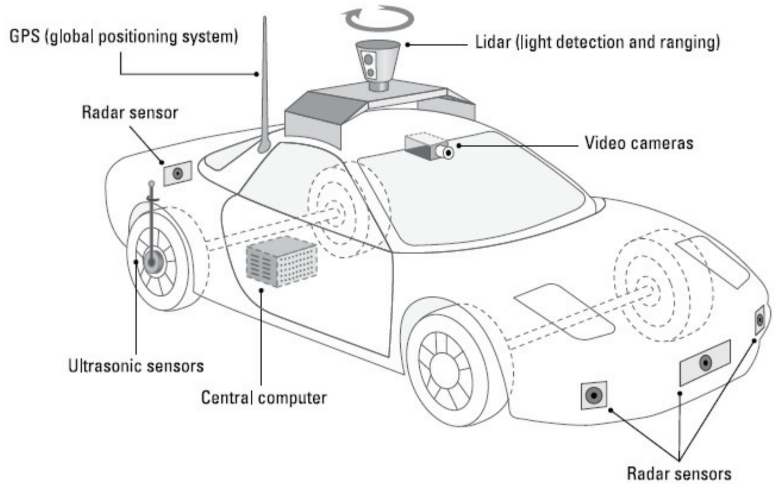

Figure 4.3. : A schematic representation of exteroceptive sensors in an SD car.

Limitations come in a number of forms. As you explore what sensors do for an SD car, you must consider cost, sensitivity to light, sensitivity to weather, noisy recording (which means that sensitivity of the sensor changes, affecting accuracy), range, and resolution. On the other hand, capabilities involve the capability to track the velocity, position, height, and distance of objects accurately, as well as the skill to detect what those objects are and how to classify them.

Camera

Cameras are passive, vision-based sensors. They can provide mono or stereo vision. Given their low cost, you can place plenty of them on the front windshield, as well as on front grilles, side mirrors, the rear door, and the rear windshield. Commonly, stereo vision cameras mimic human perception and retrieve information on the road and from nearby vehicles, whereas mono vision cameras are usually specialized in detecting traffic signs and traffic lights. The data they capture is processed by algorithms for image processing or by deep-learning neural networks to provide detection and classification information (for instance, spotting a red light or a speed-limit traffic signal). Cameras can have high resolution (they can spot small details) but are sensitive to light and weather conditions (night, fog or snow).

Lidar (LIght Detection and Ranging)

Lidar uses infrared beams (about 900 nanometer wavelength, invisible to human eyes) that can estimate the distance between the sensor and the hit object. They use a rotating swivel to project the beam around and then return estimations in the form of a cloud of collision points, which helps estimate shapes and distances. Depending on price (with higher generally meaning better), lidar can have higher resolution than radar. However, lidar is frailer and easier to get dirty than radar because it's exposed outside the car.

Radar (RAdio Detection and Ranging)

Based on radio waves that hit a target and bounce back, and whose time of flight defines distance and speed, radar can be located in the front and rear bumper, as well as on the sides of the car. Vendors have used it for years in cars to provide adaptive cruise control, blind-spot warning, collision warning, and avoidance. In contrast to other sensors that need multiple successive measurements, radar can detect an object's speed after a single ping because of the Doppler effect.

Radar comes in short-range and long-range versions, and can both create a blueprint of surroundings and be used for localization purposes. Radar is least affected by weather conditions when compared to other types of detection, especially rain or fog, and has 150 degrees of sight and 30–200 meters of range. Its main weakness is the lack of resolution (radar doesn't provide much detail) and inability to detect static objects properly.

Ultrasonic sensors

Ultrasonic sensors are similar to radar but use high-frequency sounds (ultrasounds, inaudible by humans, but audible by certain animals) instead of microwaves. The main weakness of ultrasonic sensors (used by manufacturers instead of the frailer and more costly lidars) is their short range.

Putting together what you perceive

When it comes to sensing what is around a SD car, you can rely on a host of different measurements, depending on the sensors installed on the car. Yet, each sensor has different resolution, range, and noise sensitivity, resulting in different measures for the same situation. In other words, none of them is perfect, and their sensory weaknesses sometimes hinder proper detection. Sonar and radar signals might be absorbed; lidar's rays might pass through transparent solids. In addition, it's possible to fool cameras with reflections or bad light

SD cars are here to improve our mobility, which means preserving our lives and those of others. An SD car can't be permitted to fail to detect a pedestrian who suddenly appears in front of it. For safety reasons, vendors focus much effort on sensor fusion, which combines data from different sensors to obtain a unified measurement that's better than any single measurement. Sensor fusion is most commonly the result of using Kalman filter variants (such as the Extended Kalman Filter or the even more complex Unscented Kalman Filter).

Rudolf E. Kálmán was a Hungarian electrical engineer and an inventor who immigrated to the United States during World War II. Because of his invention, which found so many applications in guidance, navigation, and vehicle control, from cars to aircraft to spacecraft, Kálmán received the National Medal of Science in 2009 from U.S. President Barack Obama.

A Kalman filter algorithm works by filtering multiple and different measurements taken over time into a single sequence of measurements that provide a real estimate (the previous measurements were inexact manifestations). It operates by first taking all the measurements of a detected object and processing them (the state prediction phase) to estimate the current object position. Then, as new measurements flow in, it uses the new results it obtains and updates the previous ones to obtain a more reliable estimate of the position and velocity of the object (themeasurement update phase), as shown in Figure 4.4.

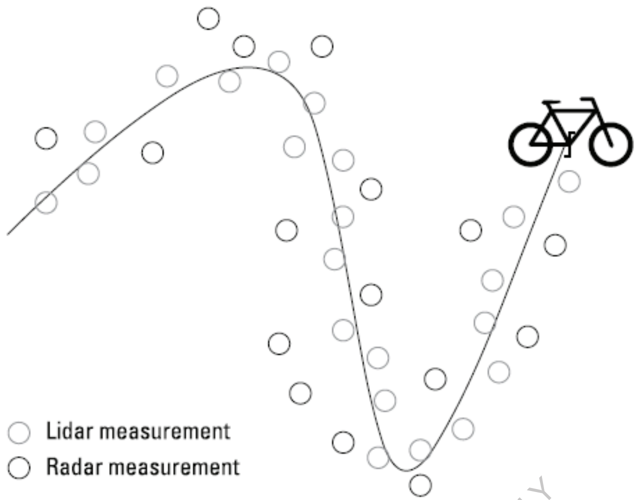

Figure4.4. : A Kalman filter estimates the trajectory of a bike by fusing radar and lidar data

In this way, an SD car can feed the algorithm the sensor measurements and use them to obtain a resulting estimate of the surrounding objects. The estimate combines all the strengths of each sensor and avoids their weaknesses. This is possible because the filter works using a more sophisticated version of probabilities and Bayes' theorem.

Understanding theNonstarter Application

AI can also fall into the trap of developing solutions to problems that don't really exist. Yes, the wonders of the solution really do look quite fancy, but unless the solution addresses a real need, no one will buy it. Technologies thrive only when they address needs that users are willing to spend money to obtain. This chapter finishes with a look at solutions to problems that don't exist.

5. USING AI WHERE IT WON'T WORK

5.1 Defining the limits of AI

When talking to Alexa, you might forget that you're talking with a machine. The machine has no idea of what you're saying, doesn't understand you as a person, and has no real desire to interact with you; it only acts as defined by the algorithms created for it and the data you provide. Even so, the results are amazing. It's easy to anthropomorphize the AI without realizing it and see it as an extension of a humanlike entity.

Creativity

You can find an endless variety of articles, sites, music, art, writings, and all sorts of supposedly creative output from an AI. The problem with AI is that it can't create anything. When you think about creativity, think about patterns of thought. For example, Beethoven had a distinct way of thinking about music. You can recognize a classic Beethoven piece even if you aren't familiar with all his works because the music has a specific pattern to it, formed by the manner in which Beethoven thought.

Creativity also implies developing a different perspective, which is essentially defining a different sort of dataset (if you insist on the mathematical point of view). An AI is limited to the data you provide. It can't create its own data; it can only create variations of existing data — the data from which it learned. The "Understanding teaching orientation" sidebar in Chapter 13 expounds on this idea of perspective. To teach an AI something new, something different, something amazing, a human must decide to provide the appropriate data orientation.

Imagination

To create is to define something real, whether it's music, art, writing, or any other activity that results in something that others can see, hear, touch, or interact with in other ways. Imagination is the abstraction of creation, and is therefore even further outside the range of AI capability. Someone can imagine things that aren't real and can never be real. Imagination is the mind wandering across fields of endeavor, playing with what might be if the rules didn't get in the way. True creativity is often the result of a successful imagination.

Just as an AI can't create new patterns of thought or develop new data without using existing sources, it must also exist within the confines of reality. Consequently, it's unlikely that anyone will ever develop an AI with imagination. Not only does imagination require creative intelligence, it also requires intrapersonal intelligence, and an AI possesses neither form of intelligence.

Imagination, like many human traits, is emotional. AI lacks emotion. In fact, when viewing what an AI can do, versus what a human can do, it often pays to ask the simple question of whether the task requires emotion.

Data deficiencies

The "Considering the Five Mistruths in Data" section of Chapter 2 tells you about data issues that an AI must overcome to perform the tasks that it's designed to do. The only problem is that an AI typically can't recognize mistruths in data with any ease unless there is an accompanying wealth of example data that lacks these mistruths, which might be harder to come by than you think. Humans, on the other hand, can often spot the mistruths with relative ease. Having seen more examples than any AI will ever see, a human can spot the mistruths through both imagination and creativity. A human can picture the mistruth in a manner that the AI can't because the AI is stuck in reality.

Mistruths are added into data in so many ways that listing them all is not even possible. Humans often add these mistruths without thinking about it. In fact, avoiding mistruths can be impossible, caused as they are by perspective, bias, and frame-of-reference at times. Because an AI can't identify all the mistruths, the data used to make decisions will always have some level of deficiency. Whether that deficiency affects the AI's capability to produce useful output depends on the kind and level of deficiency, along with the capabilities of the algorithms.

There is also the issue of speaking a hurtful truth that an AI will never be able tohandle because an AI lacks emotion. A *hurtful truth* is one in which the recipientgains nothing useful, but instead receives information that causes harm —whether emotional, physical, or intellectual. For example, a child may not knowthat one parent was unfaithful to another. Because both parents have passed on,the information isn't pertinent any longer, and it would be best to allow the child to remain in a state of bliss. However, someone comes along and ensures that thechild's memories are damaged by discussing the unfaithfulness in detail. The childdoesn't gain anything, but is most definitely hurt. An AI could cause the same sortof hurt by reviewing family information in ways that the child would neverconsider.Upon discovering the unfaithfulness through a combination of police reports, hotel records, store receipts, and other sources, the AI tells the child about the unfaithfulness, again, causing hurt by using the truth. However, in the case of the AI, the truth is presented because of a lack of emotional intelligence (empathy); the AI is unable to understand the child's need to remain in a blissful state about the parent's fidelity. Unfortunately, even when a dataset contains enough correct and truthful information for an AI to produce a usable result, the result can prove more hurtful than helpful.

5.2 Applying AI incorrectly

The limits of AI define the realm of possibility for applying AI correctly. However, even within this realm, you can obtain an unexpected or unhelpful output. For example, you could provide an AI with various inputs and then ask for a probability of certain events occurring based on those inputs. When sufficient data is available, the AI can produce a result that matches the mathematical basis of the input data. However, the AI can't produce new data, create solutions based on that data, imagine new ways of working with that day, or provide ideas for implementing a solution. All these activities reside within the human realm. All you should expect is a probability prediction.

Another issue is whether the dataset contains any sort of opinion, which is far more prevalent than you might think. An opinion differs from a fact in that the fact is completely provable and everyone agrees that a fact it truthful (at least, everyone with an open mind). Opinions occur when you don't have enough scientific fact to back up the data. In addition, opinions occur when emotion is involved. Even when faced with conclusive proof to the contrary, some humans would rather rely on opinion than fact. The opinion makes us feel comfortable; the fact doesn't. AI will nearly always fail when opinion is involved. Even with the best algorithm available someone will be dissatisfied with the output.

5.3 Entering a world of unrealistic expectations

The previous sections of the chapter discuss how expecting an AI to perform certain tasks or applying it in less than concrete situations will cause problems.

Unfortunately, humans don't seem to get the idea that the sort of tasks that many of us think an AI can perform will never come about. These unrealistic expectations have many sources, including

»»**Media:** Books, movies, and other forms of media all seek to obtain an emotional response from us. However, that emotional response is the very source of unrealistic expectations. We imagine that an AI can do something, but it truly can't do those things in the real world

»**Anthropomorphization:** Along with the emotions that media generates, humans also tend to form attachments to everything. People often name their cars, talk to them, and wonder if they're feeling bad when they break down. An AI can't feel, can't understand, can't communicate (really), can't do anything other than crunch numbers — lots and lots of numbers. When the expectation is that the AI will suddenly develop feelings and act human, the result is doomed to failure.

Deficient technology: In many places in this book, you find that a problem wasn't solvable at a certain time because of a lack of technology. It isn't realistic to ask an AI to solve a problem when the technology is insufficient. For example, the lack of sensors and processing power would have made creating a self-driving car in the 1960s impossible, yet advances in technology have made such an endeavor possible today.

5.4 Considering the Effects of AI Winters

AI winters occur when scientists and others make promises about the benefits of AI that don't come to fruition within an expected time frame, causing funding for AI to dry up and research to continue at only a glacial pace. Since 1956, the world has seen two AI winters. (Right now, the world is in its third AI summer.) The following sections discuss the causes, effects, and results of AI winter in more detail.

Understanding the AI winter

It's hard to say precisely when AI began. After all, even the ancient Greeks dreamed of creating mechanical men, such as those presented in the Greek myths about Hephaestus and Pygmalion's Galatea, and we can assume that these mechanical men would have some sort of intelligence. Consequently, one could argue that the first AI winter actually occurred sometime between the fall of the Roman empire and the time in the middle ages when people dreamed of an alchemical way of placing the mind into matter, such as JābiribnHayyān'sTakwin, Paracelsus' homunculus, and Rabbi Judah Loew's Golem. However, these efforts are unfounded

stories and not of the scientific sort that would appear later in 1956 with the founding of government-funded artificial intelligence research at Dartmouth College. An AI winter occurs, then, when funding for AI dwindles. The use of the word *winter* is appropriate because like a tree in winter, AI didn't stop growing altogether. When you view the rings of a tree, you see that the tree does continue to grow in winter — just not very fast. Likewise, during the AI winters from 1974 to 1980 and again from 1987 to 1993, AI did continue to grow, but at a glacial pace.

Defining the causes of the AI winter

The cause of an AI winter could easily be summarized as resulting from outlandish promises that are impossible to keep. At the outset of the efforts at Dartmouth College in 1956, the soon-to-be leaders of AI research predicted that a computer as intelligent as a human would take no more than a generation. Sixty-plus years later, computers still aren't nearly as smart as humans. In fact, if you've read previous chapters, you know that computers are unlikely to ever be as smart as humans, at least not in every kind of intelligence (and by now have exceeded human capability only in a very few kinds).

Part of the problem with overpromising capabilities is that early proponents of AI believed that all human thought could be formalized as algorithms. In fact, this idea goes back to the Chinese, Indian, and Greek philosophers. Only some components of human intelligence be formalized. In fact, the best possible outcome is that human mathematical and logical reasoning could be mechanized. Even so, in the 1920s and 1930s, David Hilbert challenged mathematicians to prove that all mathematical reasoning can be formalized. The answer to this challenge came from Gödel's incompleteness proof, Turing'smachine, and Church's Lambda calculus. Two outcomes emerged: Formalizing *all*mathematical reasoning isn't possible; and in the areas in which formalization is possible, you can also mechanize the reasoning, which is the basis of AI.

The end came as sort of an economic bubble. The expert systems proved brittle, even when run on specialized computer systems. The specialized computer systems ended up as economic sinkholes that newer, common computer systems could easily replace at a significantly reduced cost. In fact, the Japanese Fifth Generation Computer project was also a fatality of this economic bubble. It proved extremely expensive to build and maintain.

5.5 Rebuilding expectations with new goals

An AI winter does not necessarily prove devastating. Quite the contrary: Suchtimes can be viewed as an opportunity to stand back and think about the variousissues that came up during the rush to develop something amazing. Two majorareas of thought benefitted during the first AI winter (along with minor benefitsto other areas of thought)

Logical programming: This area of thought involves presenting a set of sentences in logical form (executed as an application) that expresses facts and rules about a particular problem domain. Examples of programming languages that use this particular paradigm are Prolog, Answer Set Programming (ASP), and Datalog. This is a form of rule-based programming, which is the underlying technology used for expert systems.

Common-sense reasoning: This area of thought uses a method of simulating the human ability to predict the outcome of an event sequence based on the properties, purpose, intentions, and behavior of a particular object. Commonsense reasoning is an essential component in AI because it affects a wide variety of disciplines, including computer vision, robotic manipulation, taxonomic reasoning, action and change, temporal reasoning, and qualitative reasoning.

The second AI winter brought additional changes that have served to bring AI intothe focus that it has today. These changes included

Using common hardware: At one point, expert systems and other uses of AI relied on specialized hardware. The reason is that common hardware didn't provide the necessary computing power or memory. However, these custom systems proved expensive to maintain, hard to program, and extremely brittle when faced with unusual situations. Common hardware is general purpose in nature and is less prone to issues of having a solution that's attempting to find a problem (see the upcoming "Creating Solutions in Search of a Problem"

Seeing a need to learn: Expert systems and other early forms of AI required special programming to meet each need, thereby making them extremely inflexible. It became evident that computers would need to be able to learn from the environment, sensors, and data provided.

Creating a flexible environment: The systems that did perform useful work between the first and second AI winters did so in a rigid manner. When the inputs didn't quite match expectations, these systems were apt to produce grotesque errors in the output. It became obvious that any new systems would need to know how to react to real-world data, which is full of errors, incomplete, and often formatted incorrectly.

Relying on new strategies: Imagine that you work for the government and have promised all sorts of amazing things based on AI, except that none of them seemed to materialize. That's the problem with the second AI winter: Various governments had tried various ways of making the promises of AI a reality. When the current strategies obviously weren't working, these same governments started looking for other ways to advance computing, some of which have produced interesting results, such as advances in robotics.

When considering AI winters and the resulting renewal of AI with updated ideas and objectives, an adage coined by American scientist and futurist, Roy Charles Amara (also known as Amara's law) is worth remembering: "We tend to overestimate the effect of a technology in the short run and underestimate the effect in the long run." After all the hype and disillusionment, there is always a time when people can't perceive the long-term impact of a new technology clearly andunderstand the revolutions it brings about with it. As a technology, AI is here to stay and will change our world for better and worse, no matter how many winters it still has to face.

5.6 Creating Solutions in Search of a Problem

Two people are looking at a mass of wires, wheels, bits of metal, and odd, assorted items that appear to be junk. The first person asks the second, "What does it do?" The second answers, "What doesn't it do?" Yet, the invention that apparently does everything ends up doing nothing at all. The media is rife with examples of the solution looking for a problem. We laugh because everyone has encountered the solution that's in search of a problem before. These solutions end up as so much junk, even when they do work, because they fail to answer a pressing need. The following sections discuss the AI solution in search of a problem in more detail

Defining a gizmo

When it comes to AI, the world is full of gizmos. Some of those gizmos really are useful, but many aren't, and a few fall between these two extremes. For example, Alexa comes with many useful features, but it also comes with a hoard of gizmos that will leave you scratching your head when you try to use them. This article by John Dvorak may seem overly pessimistic, but it provides food for thought about the sorts of features that Alexa provides: https://www.pcmag.com/commentary/354629/just-say-no-to-amazons-echo-show.

An *AI gizmo* is any application that seems on first glance to do something interesting,but ultimately proves unable to perform useful tasks. Here are some of the common aspects to look for when determining whether something is a gizmo.

(The first letter of the each bullet in the list spells the acronym CREEP, meaning,don't create a creepy AI application):

»**Cost effective:** Before anyone decides to buy into an AI application, it must prove to cost the same or less than existing solutions. Everyone is looking for a deal. Paying more for a similar benefit will simply not attract attention.

»»**Reproducible:** The results of an AI application must be reproducible, even when the circumstances of performing the task change. In contrast to procedural solutions to a problem, people expect an AI to adapt — to learn from doing, which means that the bar is set higher on providing reproducible results.

»»**Efficient:** When an AI solution suddenly consumes huge amounts of resources of any sort, users look elsewhere. Businesses, especially, have become extremely focused on performing tasks with the fewest possible resources.

»»**Effective:** Simply providing a practical benefit that's cost effective and efficient isn't enough; an AI must also provide a solution that fully addresses a need. Effective solutions enable someone to allow the automation to perform the task without having to constantly recheck the results or prop the automation up.

»»**Practical:** A useful application must provide a practical benefit. The benefit must be something that the end user requires, such as access to a road map or reminders to take medication

5.7 Avoiding the infomercial

Bedazzling potential users of your AI application is a sure sign that the application will fail. Oddly enough, the applications that succeed with the greatest ease are those whose purpose and intent are obvious from the outset. A voice recognition application is obvious: You talk, and the computer does something useful in exchange. You don't need to sell anyone on the idea that voice recognition software is useful. This book is filled with a number of these truly useful applications, none of which require the infomercial approach of the hard sell. If people startasking what something does, it's time to rethink the project.

Understanding when humans do it better

It is all about keeping humans in the loop while making use of AI. You've seen sections about things we do better than AI, when an AI can master them at all. Anything that requires imagination, creativity, the discernment of truth, the handling of opinion, or the creation of an idea is best left to humans. Oddly enough, the limits of AI leave a lot of places for humans to go, many of which aren't even possible today because humans are overly engaged in repetitive, boring tasks that an AI could easily do.

Look for a future in which AI acts as an assistant to humans. In fact, you'll see this use of AI more and more as time goes on. The best AI applications will be those that look to assist, rather than replace, humans. Yes, it's true that robots will replace humans in hazardous conditions, but humans will need to make decisions as to how to avoid making those situations worse, which means having a human at a safe location to direct the robot. It's a hand-in-hand collaboration between technology and humans.

5.8 Ten Ways in Which AIHas Failed

Any comprehensive book on AI must consider the ways in which AI has failed to meet expectations. The book discusses this issue in part in other chapters, giving the historical view of the AI winters. However, even with those discussions, you might not grasp that AI hasn't just failed to meet expectations set by overly enthusiastic proponents; it has failed to meet specific needs and basic requirements. This chapter is about the failures that will keep AI from excelling and performing the tasks we need it to do to fully achieve the successes described in other chapters. AI is currently an evolving technology that is partially successful at best.

Interpreting, not analyzing

As stated many times throughout the book, an AI uses algorithms to manipulate incoming data and produce an output. The emphasis is on performing an analysis of the data. However, a human controls the direction of that analysis and must then interpret the results. For example, an

AI can perform an analysis of an x-ray showing a potential cancer tumor. The resulting output may emphasize a portion of the x-ray containing a tumor so that the doctor can see it. The doctor might not be able to see the tumor otherwise, so the AI undoubtedly provides an importantservice. Even so, a doctor must still review the result and determine whether the x-ray does indeed show cancer. As described in several sections of the book, especially with self-driving cars in Chapter 14, an AI is easily fooled at times when even a small artifact appears in the wrong place. Consequently, even though the AI is incredibly helpful in giving the doctor the ability to see something that isn't apparent to the human eye, the AI also isn't trustworthy enough to make any sort of a decision.

5.9 Considering Human Behavior

Interpretation also implies the ability to see beyond the data. It's not the ability create new data, but to understand that the data may indicate something other than what is apparent. For example, humans can often tell that data is fake or falsified, even though the data itself presents no evidence to indicate these problems. An AI accepts the data as both real and true, while a human knows that it's neither real nor true. Formalizing precisely how humans achieve this goal is currently impossible because humans don't actually understand it.

Considering consequences

An AI can analyze data, but it can't make moral or ethical judgements. If you ask an AI to make a choice, it will always choose the option with the highest probability of success unless you provide some sort of randomizing function as well. The AI will make this choice regardless of the outcome. The "SD cars and the trolley problem" sidebar in Chapter 14 expresses this problem quite clearly. When faced with a choice between allowing either the occupants of a car or pedestrians to die when such a choice is necessary, the AI must have human instructions available to it to make the decision. The AI isn't capable of considering consequences and is therefore ineligible to be part of the decision-making process.

Discovering

An AI can interpolate existing knowledge, but it can't extrapolate existing knowledge to create new knowledge. When an AI encounters a new situation, it usually tries to resolve it as an existing piece of knowledge, rather than accept that it's something new. In fact, an AI has no method for creating anything new, or seeing it as something unique. These are human expressions that help us discover new things, work with them, devise methods for interacting with them, and create new methods for using them to perform new tasks or augment existing tasks. The following sections describe how an AI's inability to make discoveries keeps it fromfulfilling the expectations that humans have of it.

Devising new data from old

One of the more common tasks that people perform is *extrapolation* of data; for example, given A, what is B? Humans use existing knowledge to create new knowledge of a different sort. By knowing one piece of knowledge, a human can make a leap to a new piece of knowledge, outside the domain of the original knowledge, with a high probability of success. Humans make these leaps so often that they become second nature and intuitive in the extreme. Even children can make such predictions with a high rate of success.

The best that an AI will ever do is to *interpolate* data for example, given A and B, is C somewhere in between? The capability to successfully interpolate data means that an AI can extend a pattern, but it can't create new data. However, sometimes developers can mislead people into thinking that the data is new by using clever programming techniques. The presence of C looks new when it truly isn't. The lack of new data can produce conditions that make the AI seem to solve a problem, but it doesn't. The problem requires a new solution, not the interpolation of existing solutions.

Empathizing

Computers don't feel anything. That's not necessarily a negative, but this chapter views it as a negative. Without the ability to feel, a computer can't see things from the perspective of a human. It doesn't understand being happy or sad, so it can't react to these emotions unless a program creates a method for it to analyze facial expressions and other indicators, and then act appropriately. Even so, such a reaction is a canned response and prone to error. Think about how many decisions you make based on emotional need rather than outright fact. The followingSections discuss how the lack of empathy on the part of an AI keeps it from interactingwith humans appropriately in many cases.

Walking in someone's shoes

The idea of *walking in some else's shoes* means to view things from another person's perspective and feel similar to how the other person feels. No one truly feels precisely the same as someone else, but through empathy, people can get close. This form of empathy requires strong intrapersonal intelligence as a starting point, which an AI will never have unless it develops a sense of self (the *singularity* as described at https://www.technologyreview.com/s/425733/paul-allen-thesingularity-isnt-near/). In addition, the AI would need to be able to feel, something that is currently not possible, and the AI would need to be open to sharing feelings with some other entity (generally a human, today), which is also impossible. The current state of AI technology prohibits an AI from feeling or understanding any sort of emotion, which makes empathy impossible.

Of course, the question is why empathy is so important. Without the ability to feel the same as someone else, an AI can't develop the motivation to perform certain tasks. You could order the AI to perform the task, but there the AI would have no motivation on its own. Consequently, the AI would never perform certain tasks, even though the performance of such tasks is a requirement to build skills and knowledge required to achieve human-like intelligence.

Developing true relationships

An AI builds a picture of you through the data it collects. It then creates patterns from this data and, using specific algorithms, develops output that makes it seem to know you — at least as an acquaintance. However, because the AI doesn't feel, it can't appreciate you as a person. It can serve you, should you order it to do so and assuming that the task is within its list of functions, but it can't have any feeling for you.

When dealing with a relationship, people have to consider both intellectual attachment and feelings. The intellectual attachment often comes from a shared benefit between two entities. Unfortunately, no shared benefit exists between an AI and a human (or any other entity, for that matter). The AI simply processes data using a particular algorithm. Something can't claim to love something else if an order forces it to make the proclamation. Emotional attachment must carry with it the risk of rejection, which implies self-awareness.

Changing perspective

Humans can sometimes change an opinion based on something other than the facts. Even though the odds would say that a particular course of action is prudent, an emotional need makes another course of action preferable. An AI has no preferences. It therefore can't choose another course of action for any reason other than a change in the probabilities, a *constraint* (a rule forcing it to make the change), or a requirement to provide random output.

Making leaps of faith

Faith is the belief in something as being true without having proven fact to back up such belief. In many cases, faith takes the form of *trust,* which is the belief in the sincerity of another person without any proof that the other person is trustworthy. An AI can't exhibit either faith or trust, which is part of the reason that it can't extrapolate knowledge. The act of extrapolation often relies on a hunch, based on faith, that something is true, despite a lack of any sort of data to support the hunch. Because an AI lacks this ability, it can't exhibit insight — a necessary requirement for human-like thought patterns.

References

1. Ahn, Luis von. June 2006. Games with a purpose. IEEE Computer Magazine, 96–98.
2. Anusuya, M. A. and S. K. Katti. 2009. Speech recognition by machine: a review.International Journal of Computer Science and Information Security (IJCSIS)no. 6(3):181–205.
3. Bajaj, Vikas. April 25, 2010. Spammers pay others to answer security tests. New York Times.
4. Bergmair, Richard. December 2004. Natural Language Steganography and an ‘‘AI-Complete’’ Security Primitive. In 21st Chaos Communication Congress, Berlin.
5. Bishop, Mark. 2009. Why computers can't feel pain. Minds and Machines 19(4):507–516.
6. Legg, Shane and Marcus Hutter. December 2007. Universal intelligence: a definition of machine intelligence. Minds and Machines 17(4):391–444.
7. Horvitz, E. 2007. Reflections on challenges and promises of mixed-initiative interaction. AI Magazine—Special Issue on Mixed-Initiative Assistants 28(2): 11–18.
8. Hirschman, L., and R Gaizauskas. 2001. Natural language question answering. The view from here. Natural Language Engineering 7(4):275–300.
9. Thirunavukkarasu K, Ajay S. Singh, Md Irfan and Abhishek Chowdhury, "Prediction of Liver Disease using Classification Algorithms", International Conference on Computing Communication and Automation (I.C.C.C.A.), IEEE, 2018.
10. Al-Masni MA, Al-antari MA, Choi M-T, Han S-M, Kim T-S (2018) Skin lesion segmentation in dermoscopy images via deep full resolution convolutional networks. Comput Methods Programs Biomed 162:221–231.
11. Balch CM, Gershenwald JE, Soong S, Thompson JF, Atkins MB, Byrd DR, Buzaid AC et al (2009) Final version of 2009 AJCC melanoma staging and classifcation. J Clin Oncol 27(36):6199.
12. Unver HM, Ayan E (2019) Skin lesion segmentation in dermoscopic images with combination of YOLO and grabcut algorithm. Diagnostics 9(3):72.
13. Harangi B. Skin lesion classifcation with ensembles of deep convolu-tional neural networks. J Biomed Inform. 2018;86:25–32. https://doi.org/10.1016/j.jbi.2018.08.006.

14. Nida N, Irtaza A, Javed A, Yousaf MH, Mahmood MT. Melanoma lesion detection and segmentation using deep region based convolutional neural network and fuzzy C-means clustering. Int J Med Inform. 2019. https://doi.org/10.1016/j.ijmedinf.2019.01.005.
15. Goyal M, Yap MH. Automatic lesion boundary segmentation in dermo‑scopic images with ensemble deep learning methods. 2019. http://arxiv.org/abs/1902.00809. Accessed 16 July 2019.
16. Jafari MH, Nasr-Esfahani E, Karimi N, Soroushmehr SMR, Samavi S, Najarian K. Extraction of skin lesions from non-dermoscopic images using deep learning. Int J Comput Assist Radiol Surg. 2016. https://doi.org/10.1007/s11548-017-1567-8.
17. Velez DR, et al. A balanced accuracy function for epistasis modeling in imbalanced datasets using multifactor dimensionality reduction. Genet Epidemiol. 2007;31(4):306–5.https://doi.org/10.1002/gepi.20211.
18. Fawcett T. An introduction to ROC analysis. Pattern Recognit Lett. 2006. https://doi.org/10.1016/j.patrec.2005.10.010.
19. Ambad, Pravin S., and A. S. Shirat, "A Image Analysis System to Detect Skin Diseases," IOSR Journal of VLSI and Signal Processing (IOSR-JVSP), Volume 6, Issue 5, Ver. I (Sep. - Oct.2016), PP 17-25.
20. Golmei Shaheamlung, Harshpreet Kaur and Mandeep Kaur, "Survey on machine learning techniques for the diagnosis of liver disease", International Conference on Intelligent Engineering and Management, IEEE, 2020.
21. AJAY SHRESTHA and AUSIF MAHMOOD, "Review of Deep Learning Algorithms and Architectures", IEEE, 2019.
22. Yuanyuan Jia, Zhiren Tan and Junxing Zhang, "DKDR: An Approach of Knowledge Graph and Deep Reinforcement Learning for Disease Diagnosis", IEEE Intl Conf on Parallel & Distributed Processing with Applications, Big Data & Cloud Computing, Sustainable Computing & Communications, Social Computing & Networking, 2019.
23. Meherwar Fatima and Maruf Pasha, "Survey of Machine Learning Algorithms for Disease
24. Diagnostic", Journal of Intelligent Learning Systems and Applications, 2017.
25. Animesh Urgiriye and Rupali Bhartiya," Review of Machine Learning Algorithm on Cancer Data Set", International Journal of Scientific Research & Engineering Trends, Volume 6, Issue 6, Nov-Dec-2020.

26. Adnan Mohsin Abdulazeez, Maryam Ameen Sulaiman, Diyar Qader Zeebaree, "Evaluating Data Mining Classification Methods Performance in Internet of Things Applications", Journal Of Soft Computing And Data Mining, Vol.1, NO. 2, pp. 11-25, 2020.
27. Aji Prasetya Wibawa and et al, “Naive Bayes Classifier for Journal Quartile Classification”, IJES – Vol. 7, No. 2, 2019.


More Books!
yes
I want morebooks!