PEARSON BUSINESS ANALYTICS SERIES

SECOND EDITION

ARTIFICIAL INTELLIGENCE *FOR* BUSINESS

WHAT YOU NEED TO KNOW ABOUT MACHINE LEARNING AND NEURAL NETWORKS

DOUG ROSE

FREE SAMPLE CHAPTER

SHARE WITH OTHERS

Artificial Intelligence for Business

This page intentionally left blank

Artificial Intelligence for Business

What You Need to Know about Machine Learning and Neural Networks

Doug Rose

Addison-Wesley

Boston • Columbus • New York • San Francisco • Amsterdam • Cape Town
Dubai • London • Madrid • Milan • Munich • Paris • Montreal • Toronto • Delhi • Mexico City
São Paulo • Sydney • Hong Kong • Seoul • Singapore • Taipei • Tokyo

Editor-in-Chief: Mark L. Taub

Acquisitions Editor: Malobika Chakraborty

Development Editor: Chris Zahn

Managing Editor: Sandra Schroeder

Senior Project Editor: Lori Lyons

Production Manager: Aswini Kumar/codeMantra

Copy Editor: Gill Editorial Services

Indexer: Erika Millen

Proofreader: Betty Pessagno

Cover Designer: Chuti Prasertsith

Compositor: codeMantra

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at corpsales@pearsoned.com or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact intlcs@pearson.com.

Visit us on the Web: informit.com/aw

Library of Congress Control Number: 2020947363

ISBN-13: 978-0-13-655661-9

ISBN-10: 0-13-655661-2

ScoutAutomatedPrintCode

For Jelena and Leo

This page intentionally left blank

Contents at a Glance

Contents

Foreword

Approximately two decades ago, when I first started to explore artificial intelligence (AI) and machine learning (ML), these technologies were, for the most part, confined to academic institutions. In the past few years, we have seen a major leap in the field as AI and ML have begun to be used successfully in real-world applications, such as driverless cars, targeted advertising, prediction of rare diseases, and so on. Three major factors are behind this development: the availability of big data, improvements in processing power, and development of advanced ML algorithms. Today even a mobile phone is capable of running sophisticated ML algorithms in real time. These technological achievements introduce considerable opportunities for researchers and nonexperts to learn, use, and deploy customized ML solutions for their specific needs. Flexible software libraries are being made available for public use, and the models can be built in the cloud, which makes them an even more accessible solution.

Unfortunately, the simplicity and availability of software libraries hide the complexities of ML solutions. On the other extreme are math-heavy books that reveal and revel in the complexities of ML but in a language that only experts and researchers can understand. What have been lacking are books and other credible content that tell the wonderful story of AI and ML in an easy-to-understand format. This book, *Artificial Intelligence for Business: What You Need to Know about Machine Learning and Neural Networks*, fills that gap, serving as a bridge for nontechnical executives to understand AI and ML and use these technologies to solve their business problems.

I have published many research papers in international journals and conferences in the field of AI and ML. I thought this book would be one of those reviews that I do as part of my academic service. However, after reading Chapter 1, "What Is Artificial Intelligence?", I could tell that the tone was different. It was clear that Doug knows his audience, and he wants to tell a simple yet compelling story. His goal was to attract everyone to this fascinating world of AI and ML without overwhelming or frightening them with the mathematical jargon. As data science becomes more pervasive than ever, we need such efforts to introduce the field of AI and ML to everyone and expose them to the possibilities they create in our world of ever-accelerating change.

The book is divided into four major parts: Thinking Machines: An Overview of Artificial Intelligence, Machine Learning, Artificial Neural Networks, and Putting Artificial Intelligence to Work. In Part I, Doug dedicates a chapter to explaining AI fundamentals, beginning with its history. A basic understanding of the evolution of AI and what AI can and cannot do is essential for grasping the possibilities and limitations. Doug goes on to explain the fascinating notion of strong and weak AI, the power of combining AI with big data, and the ensuing challenges. He spends considerable time explaining fundamental concepts of the field, including expert systems, data mining, supervised and unsupervised learning, backpropagation of errors, and regression analysis. This part of the book also emphasizes the major applications of AI, such as intelligent robots, natural language processing (NLP), and the Internet of Things (IoT). By the end of this part, readers' minds already will be busy thinking up new ways to harness the power of AI and ML to enhance their business.

Part II dives deeper into the core concepts of ML, illustrating the concepts with numerous examples that enable readers to relate ML to the specific business problems they want to solve. This part introduces readers to specialized ML disciplines, such as semi-supervised learning, reinforcement learning, ensemble learning, and popular algorithms, without exposing them to heavy mathematics but through the use of examples and analogies. Exposing beginners to different models of ML is key to expanding their lateral understanding of the field and is something this part does well. In particular, the archery analogy used to describe the bias-variance trade-off is an incredible explanation of this challenging concept.

Part III is dedicated to neural networks. This is the most challenging part because literature on neural networks without major mathematical equations is scarce. However, Doug starts off nicely by providing a brain analogy while differentiating the brain from neural networks. He then explains the perceptron and activation functions and how they work by providing simple calculations to make readers understand the processing power of the fundamental unit of a neural network. The current research in deep learning has given rise to complex and large neural networks. However, by giving readers a fundamental understanding of a single perceptron's form and function, Doug prepares readers to tackle the more challenging concepts surrounding modern deep neural networks, which may have thousands of such entities.

Because beginners often struggle to understand concepts such as backpropagation, gradient descent optimization, and the cost function, Doug presents these ideas succinctly with examples from everyday life in Part IV. He connects neural networks with the ML concepts already learned in Part II. This approach helps readers make sense of different ML concepts and their relationship to each other and to neural networks in solving different challenging problems. This part ends by identifying key challenges, providing guidance on choosing the right type of data and tools, and explaining the importance of taking an exploratory approach.

I highly recommend this book to beginners in the field of AI and ML. Even if you have no background in mathematics or statistics, you will find the book easy to understand. This book is also a good read for many ML practitioners and data scientists, who may want a refresher on certain key concepts. I am sure, after reading this book, you will be filled with new ideas and be eager to further explore the fascinating field of AI and ML. This is the purpose of this book, and Doug met that goal. I wish Doug and this book all the best. It is an honor and a privilege to be associated with this book.

—Dr. Shehroz S. Khan
University of Toronto

This page intentionally left blank

Preface

For thousands of years, people have been fascinated with artificial intelligence (AI). In ancient mythology, the Greek god Hephaestus was so skilled with his hammer that he built a giant bronze mechanical man called Talos to protect Europa in Crete from pirates. Thousands of years later, in 1817, Mary Shelley grappled with AI when she wrote *Frankenstein*. Even more recently, the classic 1927 futuristic film *Metropolis* featured the robot named Maria who cared for the children and ultimately drove the city to rebellion.

While myths and science fiction have stirred the imagination to consider the potential (and potential horrors) of AI, philosophers have struggled to define the very nature of human intelligence. In the 1600s, Thomas Hobbes famously wrote that our "reason is nothing but reckoning." He concluded that we are just the sum of our memories, suggesting that maybe these memories could be coded to mechanical intelligence.

A few centuries later, in the 1960s, the philosopher Hubert Dreyfus criticized this idea of mechanical intelligence, writing several books on the subject. One of the most famous was *What Computers Can't Do: A Critique of Artificial Reason*, which was first published in 1972. His key argument was that our unconscious human instincts could never be captured with formal rules.

Movies like *The Terminator* warn us about how easy it might be to create an intelligent network like Skynet that could extinguish our species through a simple binary miscalculation. Silicon Valley celebrities have argued whether AI will be our greatest achievement or the ultimate cause of human extinction.

Mythology, science fiction, and philosophy make the topic of AI fascinating. However, when you start to look at how artificially intelligent machines do what they do, it's hard to be sure whether to scream or yawn. On the one hand, some displays of AI are extraordinary. It's amazing to take a ride in a self-driving car. On the other hand, most modern AI focuses on classification. You have a machine classifying millions of photos, videos, or audio files. That's not the kind of technology that motivates you to build an underground bunker or start smashing robots.

Certainly, AI has enormous potential. But we tend to like to judge things based on their performance, not their potential. So far, AI has performed well. The availability of massive data sets over recent years has given machines new food to find out more about us and the world in which we live. Machines are able to identify patterns in data that humans can't perceive and would probably never think to look for. But there's still an enormous gap between this level of performance and human intelligence.

There's also an enormous gap between the threats posed by AI and human nature. Certainly, AI carries practical and ethical challenges. But the first round of the challenges will be less about the ethical implications of creating sentient beings and more about our responsibilities to each other. Think of it as less like *The Terminator* and more like the 1981 cult classic *Escape from New York*. When the movie came out, unemployment among 16-to-24 year-olds males was at 84%. It imagined New York as a compassionless, lawless urban jungle that had to be converted into a prison. Our ethical obligations to each other will dog us long before any existential threats from rogue robots.

The first challenge posed by AI will almost certainly be how to support the people whose skills will be obsolete through automation. What will we do with the tens of millions of truck drivers, cabdrivers, retail workers, machine operators, and accountants? They won't all become programmers, yoga instructors, personal trainers, YouTubers, and artists.

It's much more likely that these socioeconomic challenges with automation will eventually eclipse our concerns about machines outsmarting humans. It won't be about a supercomputer taking control of a robot army and turning it against the human race. Instead, it will be about the automated burger flipper that took your nephew's job at Steak and Shake. After all, he may have needed that job to help pay for college.

You should be aware of these challenges as you start to think about the impact of AI. But this book isn't about grappling with these socioeconomic challenges. It's about opportunities. Specifically, it's about business opportunities. To find the best business opportunities, you need to better understand AI as a *tool*.

If you think about it, some of the top businesses didn't succeed because they were first to market. Apple didn't build the first music player. Google wasn't the first search engine. These companies succeeded because they understood the scope of the tools and technologies and how to apply them to current and future business needs.

This book is about getting you on that path. You'll get a high-level overview of the different technologies under the umbrella of AI. Throughout the book you'll see examples of how to apply this technology to different business opportunities. Once you better understand the tools, you'll be in a much better place to create long-term strategies for a new or existing business.

The business opportunities are too diverse (and as yet to be discovered) for a simple list. My hope is that you will understand the full scope of the technology and then be able to apply it to the opportunities in your organization or even to start a new business.

This book is organized into four parts. Part I is an overview of AI. Part II expands on this overview and deepens your understanding of machine learning. Part III goes into *neural networks*—computers that simulate the structure and the function of the human brain through the use of layers of interconnected artificial neurons. Finally, Part IV covers some common tools for using AI to help your business. This area has grown in popularity in recent years due to the increasing availability and decreasing cost of computer storage and processing and accessibility to massive data sets.

In Part I you'll look at some of the early theories that drove the design of the first intelligent machines. Most of these theories start with an attempt to understand human intelligence. What does it mean to be intelligent? Is it our ability to connect symbols to concepts? Our creativity?

You'll see the struggle that early computer scientists had with trying to create the first intelligent programs. At first, many computer scientists focused on symbolic reasoning. They figured that if they could get computers to understand our symbols, it would help them better understand our world. So they created systems that identified letters in our alphabet, digits, and different graphical representations like stop signs and question marks.

These early ideas still influence AI today. This symbolic approach gave rise to expert systems. These systems went through countless if-then statements to simulate thinking and decision-making; for example, if you see *A*, then make an "ah" sound. If you see a stop sign, then stop. Each of these decision points had to be painstakingly programmed into a computer.

In the 1990s, expert systems were the dominant form of AI. Companies used these systems to help make medical diagnoses, approve or reject loan applications, or find a good stock pick. The computer would go through long lists of if-then statements. So for a loan you might have an expert system that goes through a predefined list, such as, "If they have a credit history, then check for missed payments." "If they missed payments, then how many payments were missed over the past year?" "If they missed more than 10 payments in the past year, then reject this loan application."

As you can imagine, these lists can get pretty long. You need a human to try to imagine every possible if and then. A really complex task could result in a *combinatorial explosion*—so many different possibilities that it's nearly impossible to come up with all the different combinations.

As computer programmers encountered these limitations, they started to revisit the idea of machine learning (ML). ML has actually been around since the early 1950s. It was used to create programs that could beat a human player at checkers. These checker programs were extremely innovative. The machines could come up with their own strategies and learn from their mistakes. Even these early computers were sophisticated enough to learn how to beat a human player.

ML was a huge leap from programmed instructions and if-then statements that merely simulated the very human process of thinking and making decisions. Part II of this book takes a deeper dive into ML to reveal how it changed the rules of traditional software development.

With ML, the machine no longer needs to be explicitly programmed to complete a task; it can pour through massive data sets and create its own understanding. It can learn from the data and create its own model, one that represents the different rules to explain relationships among data and use those rules to draw conclusions and make decisions and predictions.

With ML, you might feed a machine all the data on the different inventory it takes to build a car along with blueprints for every car ever manufactured. After pouring through this data, the machine begins to understand certain things about what it means to be a car. It knows that cars need wheels, doors, and a windshield. There might be thousands of different kinds of cars, but the machine creates a model to identify all of them.

To enable machines to create these models, programmers have developed numerous advanced ML algorithms. A *machine learning algorithm* is a mathematical function that enables the machine to identify relationships among inputs and outputs. The programmer's role has shifted from one of writing explicit instructions to creating and choosing the right algorithms.

To take ML to the next level, computer scientists came up with the concept of an *artificial neural network*, the topic of Part III of this book. Artificial neural networks are patterned after the structure and function of the brain. The machine contains a web of interconnected artificial "neurons," each of which contains an ML algorithm. These neurons make decisions based on inputs from other neurons, the strength of the connections to those other neurons, and the deciding neuron's own algorithm and internal bias.

The artificial neural network was inspired by the way biological neurons work in the human brain. As humans, we learn new things and create memories based on increasing the strength of the connections between these nerve cells.

Modern artificial neural networks can create ML systems consisting of billions of these neurons. Such a complex network has tremendous power to find patterns in massive data sets. You can feed data into such a network, and it will create a model to better understand the larger patterns. For example, you could feed millions of images of dogs into your neural network and let it self-adjust and create its own model of what it means to be a dog. That model might not match how humans think of dogs. It may not identify dogs by looking at their shape and color or their ears and nose. Instead, it identifies statistical patterns of the different dots (pixels) in the images of the dogs. In a sense, the neural network develops its own understanding of "dogness." This way it can learn to correctly identify a dog even if it has never seen this particular dog before.

As you can imagine, the predictive power of neural networks has a wide variety of practical applications and enormous business potential. If you're in finance, neural networks can spot trends in the market to help you make trades. If you're in pharmaceuticals, you might have a neural network look for characteristics of existing drugs and compare them to new compounds. If you're in retail, you might look for patterns in what customers buy to figure out what they're likely to buy next.

Many large companies are already using neural networks for voice recognition, transcription, and digital personal assistants. For example, if you subscribe to Netflix, the system recommends movies and shows based on what you've watched in the past. Amazon uses neural networks to make targeted product recommendations and to power its personal digital assistants, including Alexa.

But you don't need to go that big to reap the value of neural networks. Think about the data in your organization. Then think about some of the patterns that would be valuable to see in your data. If you can quickly come up with valuable patterns, then AI is probably a good fit for your organization.

If you have no experience with AI, it is probably best to read this book from start to finish. By reading the chapters in Part I, you'll develop the fundamental understanding of ML required to tackle more complex topics. If you are more familiar with AI, then you could potentially start with Part II.

As you read, keep in mind that the ultimate purpose of this book is to get you thinking about the challenges and problems in your business or your area of expertise that AI and ML can help you overcome or solve. Think about the data you have, and imagine what you could possibly extract from that data to overcome a specific challenge, solve a specific problem, or answer a specific question. After all, without your very human ability to ask questions and imagine possibilities, AI and ML are useless. The power, as it has always been, is in the combination of our creativity and our tools.

Acknowledgments

Thanks first to all the organizations that have given me the opportunity to work with their talented teams, including Home Depot, Cox Automotive, Cardlytics, Genentech, and the Center for Data Intensive Science (CDIS).

Thanks also to the University of Chicago. The students in my classes at UChicago provided great feedback, and answers to their questions are weaved throughout this book. I also have online courses as a companion to this book on LinkedIn Learning. LinkedIn is one of the finest organizations I've ever worked with, and their highly skilled employees made creating these courses both fun and interesting. Special thanks to Steve Weiss, Yash Patel, and Scott Erickson for planning, editing, and filming these courses.

This book would not have been possible without the help of quite a few people. Joe Kraynak did a terrific job editing the book. John Haney created all the wonderfully whimsical illustrations. And Shehroz Khan provided outstanding technical feedback and excellent guidance. I'd also like to thank the editors at Pearson, including Kim Spenceley, Chris Zahn, and Karen Davis.

Special thanks also to my wonderful wife, Jelena, who is always my first (and most brutal) editor. She patiently read through my first drafts while sitting in the corner of my son's karate dojo. Finally, thank you to my son Leo for being my inspiration. Please keep writing. My greatest hope is that one day I'll simply be known as Leo Rose's dad.

This page intentionally left blank

About the Author

Doug Rose has been transforming organizations through technology, training, and process optimization for more than 25 years. He's the author of the Project Management Institute's (PMI) first major publication on the agile framework, *Leading Agile Teams*. He is also the author of *Data Science: Create Teams That Ask the Right Questions and Deliver Real Value* and *Enterprise Agility for Dummies*.

Doug has a master's degree (MS) in information management, a law degree (JD) from Syracuse University, and a BA from the University of Wisconsin–Madison. He is also a Scaled Agile Framework Program Consultant (SPC), Certified Technical Trainer (CTT+), Certified Scrum Professional (CSP-SM), Certified Scrum Master (CSM), PMI Agile Certified Professional (PMI-ACP), Project Management Professional (PMP), and Certified Developer for Apache Hadoop (CCDH).

You can attend Doug's lively and engaging business and project management courses at the University of Chicago or online through LinkedIn Learning.

Doug works through Doug Enterprises, an organization with an office in whatever city he lives. Currently that's in Atlanta, Georgia, where he spends his free time either riding a stationary recumbent bike or explaining the Marvel Universe to his son.

For more about Doug, visit his website at http://www.dougenterprises.com.

Part I

Thinking Machines: An Overview of Artificial Intelligence

This page intentionally left blank

1

What Is Artificial Intelligence?

In this chapter:

- Defining *intelligence* and *artificial intelligence*
- Tracing the early history of artificial intelligence
- Recognizing key limitations
- Differentiating strong and weak artificial intelligence

In 1955, Dartmouth professor John McCarthy coined the term *artificial intelligence* (AI for short) as part of an academic grant to assemble the first AI conference—the Dartmouth Summer Research Project on Artificial Intelligence in 1956. The goal of this conference was to get computers to behave in ways that humans would identify as *intelligent*.

At the time, computers were taking up whole floors in office buildings, yet they had less processing power than most modern smart watches. Making these computers intelligent was quite an ambitious goal, and conference participants soon bumped up against the limitations inherent in the hardware at the time. They made little progress toward creating the machine equivalent of a human brain.

The most lasting contribution from this grant was the term *artificial intelligence*. It ignited everyone's imagination and inspired journalists, writers, academics, and computer scientists to envision a futuristic world in which machines would think like humans. Had Professor McCarthy come up with a different name, this conference, in all likelihood, would have faded into memory. Thanks to McCarthy's choice of words, however, *artificial intelligence* has continued to fuel the imagination and drive progress toward creating intelligent machines.

Unfortunately, the concept of AI and the prospect of machines displacing humans in the workplace are frightening to most people. Just imagine if the first personal computers had been called *artificial employees*. Workers would have panicked as soon as the first PCs arrived at their office. Personal computers sound *personable*. Artificial employees would have threatened to take their jobs!

Likewise, the term *artificial intelligence* sends shivers down the spines of many people who rely on their intelligence for their jobs. This can include professionals such as lawyers, doctors, and analysts. They might all imagine a day when they're supplanted by computerized counterparts.

To alleviate some of the fear surrounding AI, it is important to separate the term from the technology. While the term evokes images of sentient and perhaps omniscient machines supplanting humans, the technology is more subdued. You won't see a mechanized version of the human brain any time soon. As a technology, AI is merely any system that exhibits behavior that could be interpreted as human intelligence, such as winning a game of chess against a world-renowned chess master.

What Is Intelligence?

The dictionary definition of *artificial intelligence* is the capability of a machine to imitate intelligent human behavior. Determining the meaning of *intelligence*, however, is the greater challenge. Although we all agree that intelligence has something to do with knowledge and the ability to reason, human intelligence seems to go beyond that to include consciousness or self-awareness, wisdom, emotion, sympathy, intuition, and creativity. To some, intelligence also involves spirituality—a connection to a greater force or being.

To further challenge our ability to define *intelligence* is the fact that human intelligence comes in many forms. Whereas some people are highly intelligent in the field of mathematics, others excel in art, music, politics, business, medicine, law, linguistics, and so on. Some people may excel in academics, whereas others are skilled in trades or have a higher level of emotional competence. And although people have tried to develop a single standard for measuring intelligence, such as the intelligent quotient (IQ), such standards are skewed. For example, a typical IQ test evaluates only short-term memory, analytical thinking, mathematical ability, and spatial recognition.

Without a reliable standard for measuring human intelligence, it's very difficult to point to a computer and say that it's behaving intelligently. Computers are certainly very good at performing certain tasks and may do so much better and faster than humans, but does that make them intelligent? For example, computers have been able to beat humans in chess for decades. IBM Watson beat some of the best champions in the game show *Jeopardy*. Google's DeepMind has beaten the best players in the 2500-year-old Chinese game called Go—a game so complex that there are thought to be more possible configurations of the board than there are atoms in the universe. Yet none of these computers understands the purpose of a game or has a reason to play.

As impressive as these accomplishments are, they are still just a product of a computer's special talent for *pattern matching*—extracting information from its database that enables it to answer a question or perform a task. This seems to be intelligent behavior only because a computer is excellent at that particular task. However, we rarely attribute human characteristics to other machines, such as boats that can "swim" faster or hydraulic jacks that are "stronger" and can easily lift a car above a mechanic's head.

In many ways a game is a perfect environment for a computer. It has set rules with a certain number of possibilities that can be stored in a database. When IBM's Watson played *Jeopardy*, all it needed to do was use natural language processing (NLP) to understand the question, buzz in faster than the other contestants, and apply pattern matching to find the correct answer in its database.

Early AI developers knew that computers had the potential to excel in a world of set rules and possibilities. Only a few years after the first AI conference, developers had their first version of a chess program. The program could match an opponent's move with thousands of possible counter moves and play out thousands of games to determine the potential ramifications of making a move before deciding which piece to move and where to move it, and it could do so in a matter of seconds.

AI is always more impressive when computers are on their home turf—when the rules are clear and the possibilities limited. Organizations benefiting most from AI are those that work within a well-defined space with set rules, so it's no surprise that organizations like Google fully embrace AI. Google's entire business involves pattern matching—matching users' questions with a massive database of answers. AI experts often refer to this as good old-fashioned artificial intelligence (GOFAI).

If you're thinking about incorporating AI in your business, consider what computers are really good at—pattern matching. Do you have a lot of pattern matching in your organization? Does a lot of your work have set rules and possibilities? This work will be the first to benefit from AI.

Testing Machine Intelligence

Alan Turing was an English computer scientist who famously took part in decrypting the enigma machines that the Germans used to communicate during World War II. After the war, he set his sights on early computers. In particular, he was interested in how machines might be able to think.

In a 1951 paper, he proposed a test called the *imitation game* that was based on a Victorian parlor game. In the game, a man and woman sat in one room, and their *interrogator* sat in another (Figure 1.1).

Figure 1.1 Imitation game

The interrogator would ask the man and the woman a question. Then the team would pass back their answers in a written note. It was up to the interrogator to decide if each written answer came from the man or the woman. In an added twist, the man tries to fool the interrogator, whereas the woman tries to help.

Now, to a modern ear, this game sounds dreary and misogynistic. But to Turing, this was an excellent foundation to test a machine's intelligence. He imagined an updated imitation game where the man was replaced by a machine (Figure 1.2).

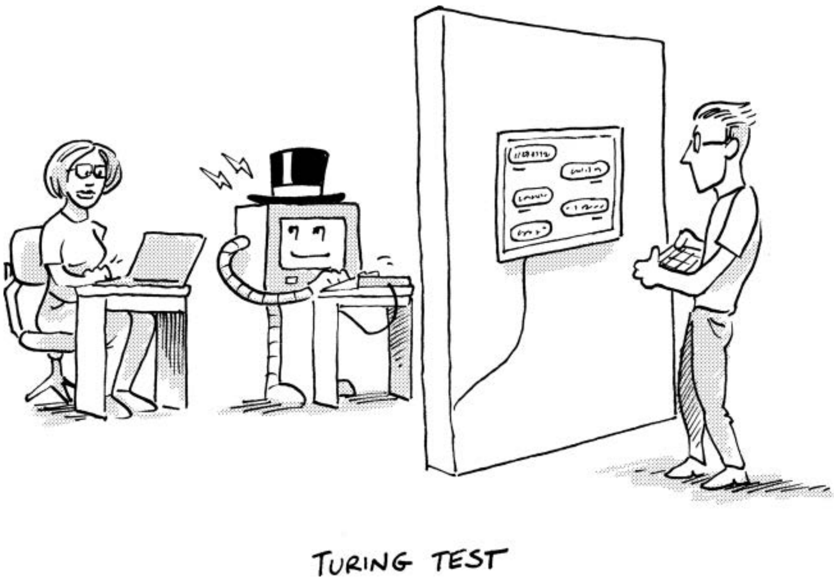

Figure 1.2 Turing test

Then the interrogator would ask both the woman and the machine a question and get back their answer in a written note. If the interrogator was just as likely to pick one or the other, then the machine must be seen as intelligent. This game was later known as the *Turing test*.

This test sparked a lot of curiosity in an "imaginable machine" even though it came out a few years before McCarthy even coined the term *artificial intelligence*. Even after nearly 70 years, this test still sounds intriguing. Imagine if you could ask a machine a question in your own natural language and get a response that is indistinguishable from that of another human?

That being said, most experts agree that the Turing test is not necessarily the best way to gauge intelligence. For one, it depends a lot on the interrogator. Some people might be easily fooled into thinking that they're talking to another person. It also assumes that AI will be similar to human intelligence. You would assume that a machine would be able to have a decent conversation before it started performing an advanced task such as searching for new drugs or accurately predicting global weather patterns.

Yet the Turing test still inspires a lot of innovation. Companies still try to create intelligent chatbots, and there are still NLP competitions that attempt to pass the test. It seems like modern machines are only a few years away from passing the Turing test. Many modern NLP applications can accurately understand the majority of your requests. Now they just need to improve their ability to respond.

Yet even if a machine can pass the test, it still seems unlikely that that same machine would be seen as intelligent. Even if your smart phone can trick you into thinking you're talking to a human, that doesn't mean that it will offer meaningful conversation.

The General Problem Solver

One of the very first attempts at AI was in 1956. Allen Newell and Herbert A. Simon (Figure 1.3) created a computer program they called the general problem solver. This program was designed to solve any problem that could be presented in the form of mathematical formulas.

Figure 1.3 Newell and Simon

Courtesy Carnegie Mellon University Libraries

One of the key parts of the general problem solver was what Newell and Simon called the physical symbol system hypothesis (PSSH). They argued that symbols were the key to general intelligence. If you could get a program to connect enough of these symbols, you would have a machine that behaved in a way similar to human intelligence.

Symbols play a big role in how we interact with the world. When we see a stop sign, we know to stop and look for traffic. When we see the word *cat*, we know that it represents a small furry feline that meows. If we see a chair, we know it's an object to sit in. When we see a sandwich, we know it's something to eat, and we may even feel hungry.

Newell and Simon argued that creating enough of these connections would make machines behave more like us. They thought a key part of human reasoning was just connecting symbols—that our language, ideas, and concepts were just broad groupings of interconnected symbols (Figure 1.4).

Figure 1.4 Interconnected symbols

But not everyone bought into this idea. In 1980, philosopher John Searle argued that merely connecting symbols could not be considered intelligence. To support his argument against the claim that computers think or at least have the potential of someday being able to think, he created an experiment called the Chinese room argument (Figure 1.5).

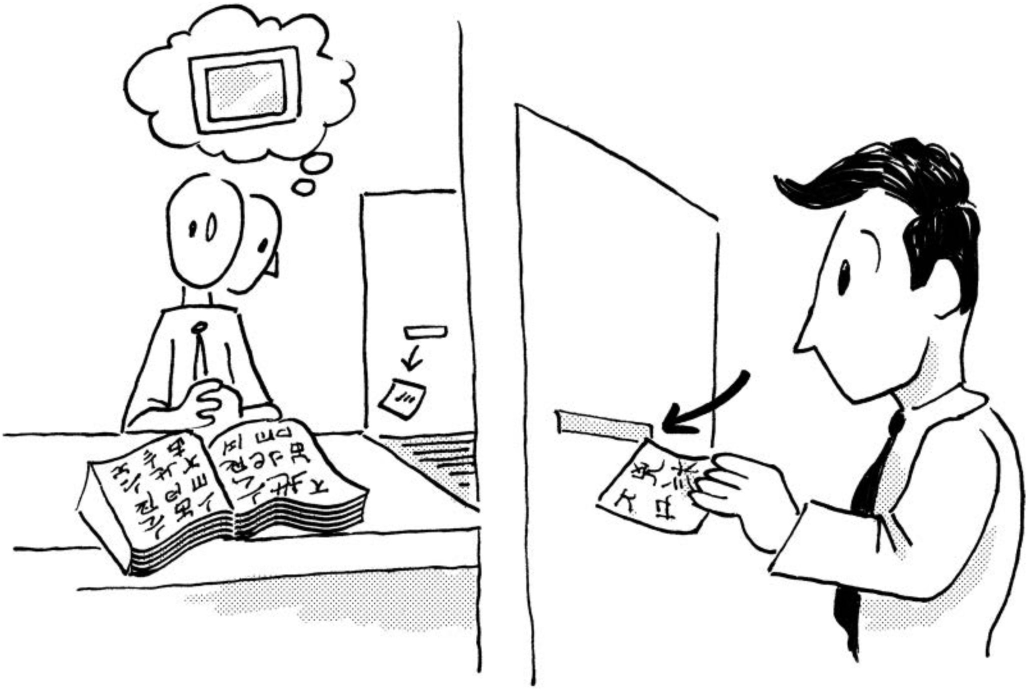

Figure 1.5 The Chinese room argument

In this experiment, imagine yourself, an English-only speaker, locked in a windowless room with a narrow slot on the door through which you can pass notes. You have a book filled with long lists of statements in Chinese, a floor covered in Chinese characters, and instructions that if you're given a certain sequence of Chinese characters you are to respond with the corresponding statement from the book.

Someone outside the room who speaks fluent Chinese writes a note on a sheet of paper and passes it to you through the slot on the door. You have no idea what it says. You go through the tedious process of looking through your book and finding the statement in response to the sequence of Chinese characters on the note. Using the characters from the floor, you paste together the statement to a sheet of paper and pass it through the slot to the person who gave you the original message.

The native Chinese speaker who passed you the note believes that the two of you are conversing and that you're intelligent. However, Searle argues that this is far from intelligence because you can't speak Chinese, and you have no understanding of the notes you're receiving or sending.

You can try a similar experiment with your smart phone. If you ask Siri or Cortana how she's feeling, she's likely to say she's feeling fine, but that doesn't mean she's feeling fine or feeling anything at all. She doesn't even understand the question. She's just matching your question to what are considered acceptable answers and choosing one.

A key drawback of matching symbols is what's referred to as the *combinatorial explosion*—the rapid growth of symbol combinations that makes matching increasingly difficult. Just imagine the variety of questions that people can ask and all the different responses to a single question. In the Chinese room example, you'd have an ever-growing book of possible inputs and outputs, which would take you longer and longer to find the correct response.

Even with these challenges, symbol matching remained the cornerstone of AI for 25 years. However, symbol matching has been unable to keep up with the growing complexity of AI applications. Early machines had trouble matching all the possibilities, and even when they could, the process took too much time.

Strong and Weak Artificial Intelligence

John Searle didn't just create the Chinese room argument. He also pointed out that you can think of AI in two different ways. He called them strong and weak AI (Figure 1.6).

- With strong AI, a machine displays all the behavior you'd expect from a person. If you're a *Star Trek* fan, this is Lieutenant Commander Data. If you prefer *Star Wars*, then this might be C3PO or R2-D2. These artificial beings have emotions, a sense of purpose, and even a sense of humor. They may learn a new language just for the joy of learning it. Some computer scientists refer to strong AI as general AI—a broad intelligence that doesn't apply only to one narrow task.

- Weak (or narrow) AI is confined to a narrow task, such as product recommendations on Amazon and Google in response to the keywords a user enters. A weak AI program doesn't engage in conversation, recognize emotion, or learn for the sake of learning; it merely does whatever job it was designed to do.

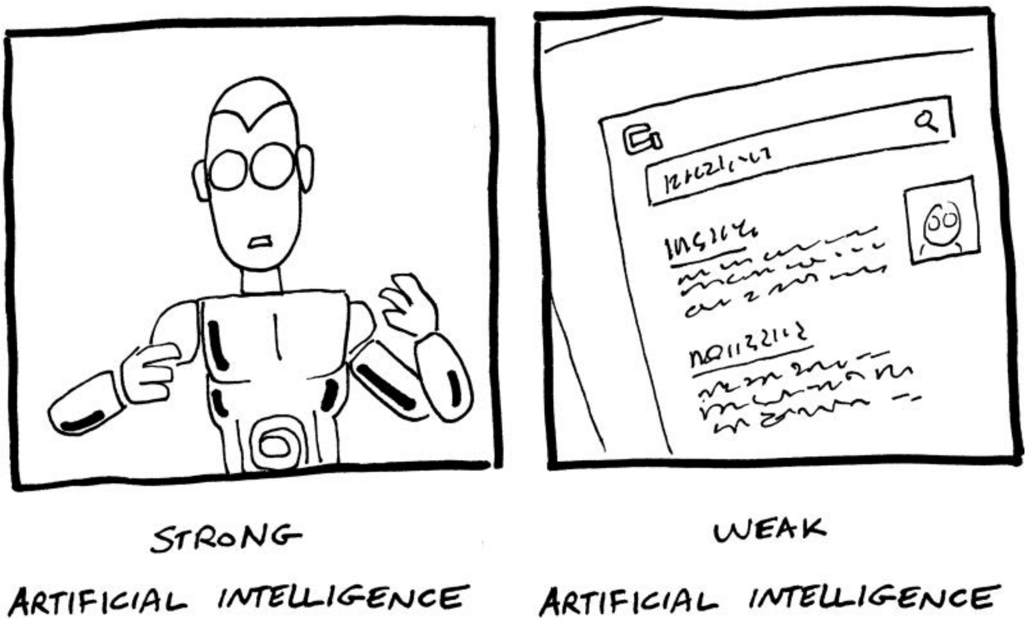

Figure 1.6 Strong and weak AI

Most AI experts believe that we're just starting down the path of weak AI—using AI to answer factual questions, provide directions, manage our schedules, make recommendations based on our past choices and reactions, help us do our taxes, prevent online fraud, and so on. Many organizations already use weak AI to help with narrow tasks, such as these. Strong AI is still relegated to the world of science fiction.

You can witness weak AI at work in the latest generation of personal assistants, including Apple's Siri and Microsoft's Cortana. You can talk to them and even ask them questions. They convert spoken language into machine language and use pattern matching to answer your questions and respond to your requests. That's not much different from traditional interactions with search engines such as Google and Bing.

The difference is that Siri and Cortana behave more like human beings; they can talk. They can even book a reservation at your favorite restaurant and place calls for you.

These personal assistants don't have general AI. If they did, they'd certainly get sick of listening to your daily requests. Instead, they focus on a narrow task of listening to your input and matching it to their database.

John Searle was quick to point out that any symbolic AI should be considered weak AI. However, in the 1970s and 1980s, symbolic systems were used to create AI software that could make expert decisions. These were commonly called *expert systems*.

In an expert system, people who specialize in a given field input the patterns that the computer can match to arrive at a given conclusion. For example, in medicine, a doctor may input groupings of symptoms that match up with various diagnoses. A nurse inputs the patient's symptoms into the computer. The computer can then search its database for a matching diagnosis and present the most likely diagnosis to the patient. For example, if a patient has a cough, shortness of breath, and a slight fever, the computer may conclude that the patient probably has bronchitis. To the patient, the computer may seem to be as intelligent as a doctor, but in reality all the computer is doing is matching symptoms to possible diagnoses.

Expert systems run into the same problems as other symbolic systems; they ultimately experience combinatorial explosions. There are simply too many symptoms, diagnoses, and variables to consider when trying to diagnose an illness. Just think about all the steps a doctor must take to arrive at an accurate diagnosis—conducting a physical exam, interviewing the patient, ordering lab tests, and sometimes ruling out a long list of other illnesses that have similar symptoms. Imagine all the possible ways a patient could answer each question the doctor asks and all the various combinations of lab results.

These early expert systems also had a serious limitation—the real possibility that given certain input, the system would be unable to find a match. You have probably experienced this on various websites; you input your search phrase, and the site informs you that it found no match.

Even with these drawbacks, the symbolic approach was a key starting point for AI and is still in use today, typically with some modifications (as you'll see next).

Artificial Intelligence Planning

Early expert systems started to disappear in the late 1980s, but the symbolic approach remained. Today, you see it in what's called *artificial intelligence planning*—a branch of AI that employs strategies and action sequences to enhance the computer's ability to match symbols and patterns.

AI planning attempts to solve the problem of combinatorial explosion by using something called *heuristic reasoning*—an approach that attempts to give AI a form of common sense, thus limiting the number of patterns the program has to match at any one time. This approach is sometimes referred to as *limiting the search space*.

Imagine heuristic reasoning applied to the Chinese room experiment. You could use heuristic reasoning in an AI program to limit the possibilities of the first note. You could set it up so that the program expects a message like "Hello" or "How are you?" or "Do you speak Chinese?" to limit how far the program has to search to match the pattern.

The drawback is that if the program doesn't receive the anticipated input, it then has to search the entire database for a match as well, which requires additional processing. For example, suppose the first note asks, "Do you know how to say purple in English?" The program must first rule out the anticipated messages and then search the entire database, or what AI planners refer to as the *entire search space*.

AI planning is common with navigation systems such as Google Maps. You enter your location and your destination, and the system finds the shortest, fastest route. It still uses a symbolic approach that relies on lists, and Google must gather the data to create those lists. It does so by pulling data from numerous sources, including satellite and aerial imagery; state, city, and county maps; the US Geological Survey; its own Street View cars; and from users who contribute their own map information. All this data is carefully vetted and then stitched together to create highly detailed maps. Google Maps also extracts current traffic data from local highway authorities to help route drivers around accidents and backups.

Unlike early symbolic systems, Google Maps uses heuristic reasoning to limit its search to a certain geographical area based on the location and destination you enter, and it can provide detailed directions, such as whether to turn left or right at a given intersection without having to search through its entire database of symbols and patterns.

Even though it's considered old-fashioned AI, symbolic systems and AI planning are still used in many new projects. It performs well in systems that have predefined symbols and patterns. You can see this with driving directions, but it also works with contracts, logistics, and even video games. If you're considering a new AI project, don't be quick to dismiss the benefits of good old-fashioned AI. Newer approaches may not be the right fit.

Learning over Memorizing

There's a big difference between memorizing and learning. Imagine that you were looking at the eight images in Figure 1.7. You would probably quickly recognize that these are eight different breeds of dogs.

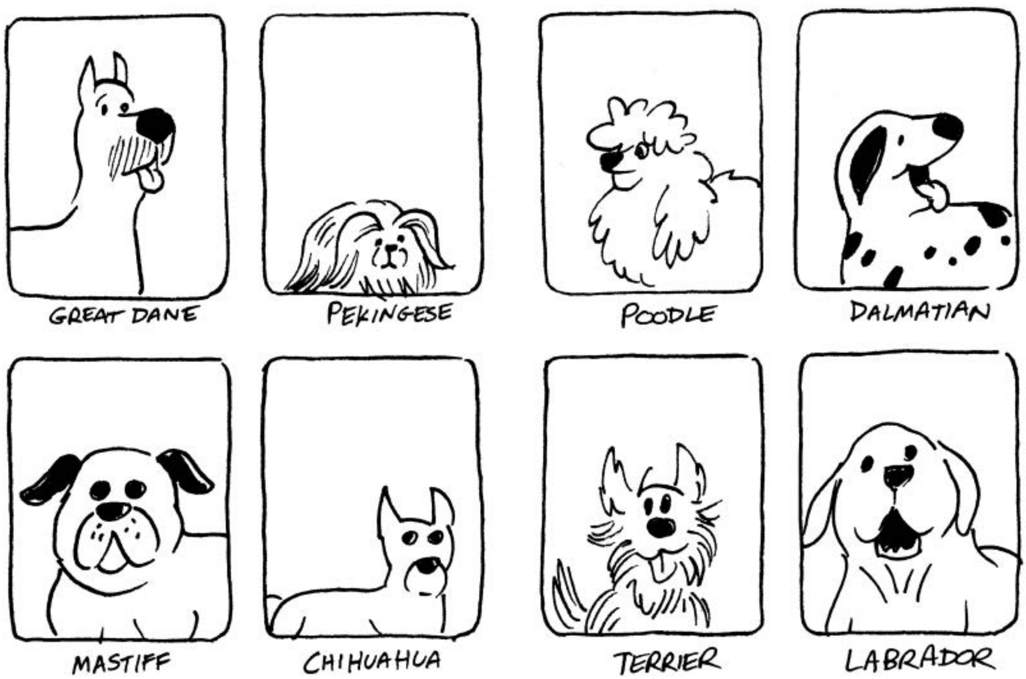

Figure 1.7 Eight breeds of dogs

Now imagine I showed you the picture in Figure 1.8.

Figure 1.8 Another dog

Would you know what it is? How would you know? You'll notice that it's not exactly the same as any of the pictures in Figure 1.7.

So that means that you've *learned* something about dogs. You might not know *what* you've learned, and you probably have trouble describing *how* you learned it's a dog, but chances are your human brain was able to make the connection.

On the other hand, computers are far better at memorization. So you could program a computer to memorize those eight dog illustrations. Then if you show the computer one of the eight, it could quickly make a connection. But it's much more difficult to have a computer learn. It could easily fumble on this ninth image as it tries to find the perfect match.

Symbolic AI is about memorizing and matching these different symbols. So a symbolic system could easily memorize the eight dog illustrations and make exact matches. But if it couldn't find an exact match, it cannot give an answer. You also could certainly create a symbolic system that translates written text from different languages. A machine could memorize millions of different typed words and phrases. The challenge is that it wouldn't really be learning a new language. It would just be a digital version of an old-style language phrasebook.

You've also seen that you can use expert systems to match your input to some preprogrammed scenario. Yet these symbolic systems will always be confined to the system's own memory. It will only know what it's programmed to know, so these systems rely on humans to *seem* intelligent. They could respond with canned phrases and memory storage. But these symbolic systems will fail when they encounter any new or unforeseen symbol (Figure 1.9).

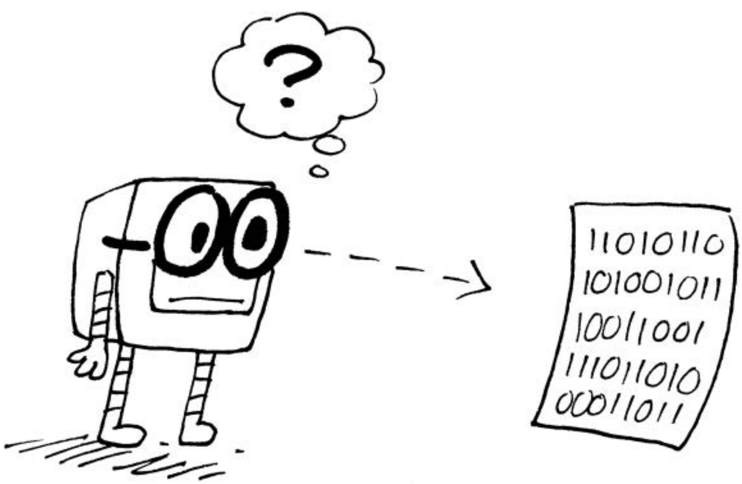

Figure 1.9 Encountering a new symbol

So the big challenge going forward is trying to make AI more *generalizable*. Instead of matching memorized symbols, newer machines will look for features and patterns to create abstract models. These models have the potential to help these machines learn. That way they can better handle new items that the machine might not have seen. If it's just predicting well on what it already knows, it's just a simple act of memorization.

Humans have evolved to be very good in both memorization and generalization. We use memorization to quickly fetch information on how to act in a familiar terrain or domain. Generalization is our human ability to use our prior knowledge to work well in unseen or unfamiliar situations. For us, survival is the key, and that is what we try to optimize all the time.

For machines, this is missing. That's why it's essential to identify the criteria a machine needs to optimize to generalize the knowledge it learns so that it can perform well in the present or future.

Chapter Takeaways

- *Artificial intelligence* is the capability of a machine to imitate intelligent human behavior.
- Early AI is pattern matching; it enables machines to imitate intelligent behavior, but it cannot be equated with human intelligence.
- Pattern matching is often enhanced with AI planning that streamlines the pattern-matching process.
- Strong AI imbues machines with human qualities, including self-awareness and emotions, whereas weak AI enables machines to perform specific tasks.
- Strong AI remains relegated to the realm of science fiction.

Index

A

B

N

O

P

Q

R

V

W

X-Y-Z