


ARTIFICIAL GENERAL INTELLIGENCE



The MIT Press Essential Knowledge Series

A complete list of books in this series can be found online at
<https://mitpress.mit.edu/books/series/mit-press-essential-knowledge-series>.

ARTIFICIAL GENERAL INTELLIGENCE

JULIAN TOGELIUS

The MIT Press

|

Cambridge, Massachusetts

|

London, England

The MIT Press
Massachusetts Institute of Technology
77 Massachusetts Avenue, Cambridge, MA 02139
mitpress.mit.edu

© 2024 Massachusetts Institute of Technology

All rights reserved. No part of this book may be used to train artificial intelligence systems or reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Chaparral Pro by New Best-set Typesetters Ltd. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Names: Togelius, Julian, author.

Title: Artificial general intelligence / Julian Togelius.

Description: Cambridge, Massachusetts : The MIT Press, [2024] | Series: The MIT Press essential knowledge series | Includes bibliographical references and index.

Identifiers: LCCN 2024000987 (print) | LCCN 2024000988 (ebook) | ISBN 9780262549349 (paperback) | ISBN 9780262380140 (epub) | ISBN 9780262380157 (pdf)

Subjects: LCSH: Artificial intelligence.

Classification: LCC Q335 .T665 2024 (print) | LCC Q335 (ebook) | DDC 006.3—dc23/eng/20240423

LC record available at <https://lcn.loc.gov/2024000987>

LC ebook record available at <https://lcn.loc.gov/2024000988>

10 9 8 7 6 5 4 3

CONTENTS

Series Foreword vii

- 1 Introduction 1
- 2 A Brief History of Superhuman AI 13
- 3 Intelligence (Natural) 33
- 4 Intelligence (Artificial) 53
- 5 Varieties of Artificial General Intelligence 69
- 6 Practical AGI Development 83
- 7 Self-Supervised Learning of Foundation Models 101
- 8 Open-Ended Learning in Virtual Worlds 127
- 9 AGI and Consciousness 145
- 10 Superintelligence and the Intelligence Explosion 155
- 11 AGI and Society 167
- 12 Conclusion 185

Acknowledgments 203

Glossary 205

Notes 209

Further Reading 221

Index 223

SERIES FOREWORD

The MIT Press Essential Knowledge series offers accessible, concise, beautifully produced pocket-size books on topics of current interest. Written by leading thinkers, the books in this series deliver expert overviews of subjects that range from the cultural and the historical to the scientific and the technical.

In today's era of instant information gratification, we have ready access to opinions, rationalizations, and superficial descriptions. Much harder to come by is the foundational knowledge that informs a principled understanding of the world. Essential Knowledge books fill that need. Synthesizing specialized subject matter for nonspecialists and engaging critical topics through fundamentals, each of these compact volumes offers readers a point of access to complex ideas.

INTRODUCTION

When you talk about artificial intelligence (AI), the term can mean many different things. It could refer to the centuries-old dream of truly intelligent machines that can reason as well as any of us. It could refer to any of the many depictions of AI in movies and books, from the Terminator to WALL-E to Commander Data from *Star Trek: The Next Generation*. It could refer to theoretical or experimental research that goes on in university labs and gets published in academic papers. Or it could refer to actual technology that we use every day, which automatically plans our driving routes, improves our photos, controls characters in our video games, does our web searches, or serves us ads based on those searches.

These different senses of AI are commonly confused in the general debate. For example, you may see a headline that reads “AI Will Take Over the World,” then look at the AI-powered autocomplete software in your phone and

conclude that something so inept is clearly not going to take over the world anytime soon. Or you may be thinking of buying a new car with AI-powered driving assistance features but worry whether it might have a (potentially murderous) mind of its own, like HAL 9000 from *2001: A Space Odyssey*. In my experience, reading an article that speculates about whether AI will make humans obsolete, and then turning my attention to my ongoing experiments where the bots I'm training run in circles and into walls instead of fighting each other, is slightly surreal. It makes you wonder whether the people who write such headlines have ever tried developing an AI system.

One source of confusion here is mistaking real software and hardware for fiction. Another source of confusion is mistaking systems that do a single task for those that can do a large variety of tasks, even those they were not trained on or built for. The fictional AI systems in *2001: A Space Odyssey*, *The Terminator*, and *WALL-E* have marvelously broad ranges of capabilities. The original Terminator, for example, can convincingly imitate the voices of people he just met, lay complex plans, do detective work, drive a motorcycle, and fight dozens of police officers with a wide range of weapons. Meanwhile, AlphaGo, the AI system that stunned the world in 2016 by beating the world's best Go players, can only play Go. It can't even play chess, unless you retrain it from scratch. It wouldn't even begin to know how to drive a motorcycle or fight the law.

This book is about *artificial general intelligence* (AGI), which is a term that is not clearly defined. Many have tried to define AGI, but the definitions are often unclear and sometimes conflict with each other. For example, Bill Gates defines AGI succinctly as “software that’s capable of learning any task or subject.”¹ John Carmack, legendary game technologist and cocreator of *Doom*, instead defines it as “a form of AI that goes beyond mimicking human intelligence to understanding things and solving problems.”² The AI researcher Murray Shanahan defines AGI as “artificial intelligence that is not specialized to carry out specific tasks, but can learn to perform as broad a range of tasks as a human.”³ It is easy to think of various systems that would be covered by one of these definitions but not by another. There are plenty more definitions to consider if one is so inclined; the AI researcher Ben Goertzel analyzes dozens of AGI definitions in a comprehensive paper from 2014, and the discussion about AGI has only intensified since then.⁴

In this book, I take *AGI* to mean software that can do a wide variety of things and solve a wide variety of problems and is significantly more capable than the AI systems we have today. We can contrast this with *narrow AI*, which refers to systems that do only one thing or a few closely related things. All, or almost all, existing AI systems are rather narrow in scope.

Another term for which we have no single generally agreed-on definition is *artificial intelligence*. A surprising

number of sometimes contradictory definitions have been proposed since the term was first used in the 1960s, partly reflecting the great variety of computational approaches that have been pursued under the heading of artificial intelligence.⁵ Some of these definitions emphasize the modeling of biological or human intelligence, others connect to notions of intelligence from psychology, yet others focus on specific computational concepts such as symbols or functions, and others still are concerned with what types of tasks an AI system could solve. In this book, I use the term *artificial intelligence* to mean two different things. The first is the quest to make machines do things that humans currently do better, specifically things that seem to require thought when humans do them. (In practice, “machines” means digital computers, at least for now.) The second meaning of *AI* is a set of algorithms, software, and insights that have been developed by researchers and programmers focusing on solving problems that seem to require intelligence. While this book is about AI as a quest to build intelligence in machines, I will often discuss individual AI technologies that have been developed.

An expression you will not see in this book, and should not see from any knowledgeable writer on the topic, is “an AI.” No such thing exists in the world today, only various pieces of software and hardware that have been developed to try to achieve various capabilities. To avoid confusion, I will instead talk about AI systems, AI tools, and AI

algorithms when I talk about the outcomes of AI research. It is important to remember that the various things we label “AI” are often quite different from each other, and a program that plays chess has little in common with a program that generates lifelike images.

The topic of the book is AGI, and given that there are no AGIs in the world yet—or, more specifically, the most general AI systems are not yet very general—one might wonder where to start. The obvious place is history—not the history of AI in general but the history of superhuman AI. The next chapter therefore digs into the six-decade history of AI to describe several examples of AI systems that do things better than humans do. The main point of the chapter is that superhuman capabilities in some problem domain, while impressive and often useful, are not necessarily a step toward general intelligence. Along the way, I will try to convey something of the rather diverse technical methods that underpin the various AI systems.

We are talking about artificial intelligence, but we have not yet clarified what we mean by “intelligence.” Chapter 3 is about intelligence of the nonartificial kind. Intelligence has been studied by philosophers, psychologists, and ethologists (those who study animal behavior), among others. We will look at a couple of key concepts in the study of intelligence in humans and animals, and how they relate to artificial general intelligence. This includes the key distinction between fluid and crystallized

A program that
plays chess has little
in common with a
program that generates
lifelike images.

intelligence, which we can use to understand the differences between different AI methods. It also includes the *g* factor, which is meant to be a measure of general intelligence in humans.

But we can draw on relevant theory from within AI research as well. Chapter 4 examines attempts to define AGI by the AI community. It also discusses some mathematical arguments for and against the possibility of AGI. These include the counterintuitive “no free lunch” theorems, which seem to say that general intelligence is impossible, but also the concept of universal intelligence, which attempts to overcome this theoretical limitation.

All these theories about intelligence and artificial intelligence can feel quite abstract. They don’t necessarily help us imagine what AGI would be like. In chapter 5, we explore some visions of what AGI could be like, with ample reference to science fiction. Science fiction stories have inspired generations of AI researchers and can help us not only think about but also differentiate between potential AI futures.

Stepping away from the abstract argument and into the world of hands-on AI research, the next three chapters discuss practical approaches to creating AGI. Chapter 6 reviews the main directions in AI research and attempts to assess whether either of them, on their own, is likely to lead to AGI. Here’s a spoiler: no.

The next two chapters examine the approaches toward AGI that I consider the most promising. The first of these,

discussed in chapter 7, is foundation models constructed through self-supervised learning. The last few years have seen amazing results in training very large neural networks on enormous amounts of data scraped from the internet. The resulting networks are sometimes shockingly capable, including performing tasks they were not explicitly trained for. OpenAI's GPT-4 was only trained to predict text but can do translation, summarization, philosophical argumentation, your homework, and many other things. Another foundation model, Stable Diffusion, can generate an enormous variety of images in response to text prompts and also compress images, retouch them, and so on. But it is also very easy to find things that these marvelous models cannot (yet) do. Will this approach scale? Can we simply keep training larger networks on more and more data and get to some form of AGI?

Chapter 8 describes another, complementary approach to AGI, namely, open-ended learning in virtual environments. In this approach, AI systems are agents that take actions in simulated worlds instead of passively learning from human-made data. The idea here is to let AI emerge or evolve, inspired by how intelligence evolved on Earth. For a long time now, scientists have studied agents that learn to act so as to maximize rewards, but while this research has shown much progress, these agents tend to learn brittle strategies that are the opposite of general. What's new with open-ended learning is that the AI sys-

tem invents its goals as it goes, and generates the problems it thinks it should try to solve. This avoids the trap of overspecialization. We have seen some encouraging results, but will such systems eventually learn skills that are relevant to us?

Now, let's assume we do manage to construct an artificial intelligence system of considerable generality. Will there be something it is like to be that system? Will it have a point of view, and will it know about itself? If so, does that give us ethical obligations toward the system? Chapter 9 tackles the tricky question of AGI consciousness. While we cannot say much with any certainty here, some interesting arguments have been proposed.

Armed with these insights, we are ready to face the most notorious idea in AGI, the main reason that so many people are not only excited by but also scared of the development of AGI. Chapter 10 is devoted to the superintelligence (or intelligence explosion) argument. In a sentence, the argument is that a sufficiently smart AI system would learn how to improve itself, and thus get even smarter, and so on, in an exponential growth of intelligence. Intellectually, we would be like ants to such a system. Does this argument check out? Do its assumptions jibe with what we know about AI?

Chapter 11 addresses the effects of AGI on society. What would a hypothetical AGI system mean for employment, and for our sense of self-worth? Would it limit our

What would a
hypothetical AGI system
mean for employment,
and for our sense of
self-worth?

ability to govern our society? Considering AGI in the context of our whole society is illuminating, as it helps us to further understand the range of things humans do and tasks that a truly general AI would need to be able to solve.

Finally, in chapter 12, I will attempt to synthesize the book into a pithy conclusion. As we reach the end of the book, it will be almost impossible for me to withhold some personal opinions. Spice level: medium.

But we're not there yet. First, let's go back to ancient digital history.

A BRIEF HISTORY OF SUPERHUMAN AI

We already have machines with superhuman intelligence, meaning that they are better than most humans on some task that would seem to require intelligence. Such superhumanly intelligent machines have existed for a long time, decades at the least. Exactly for how long depends a little on which tasks you think require intelligence. This seems to have changed over the years.

The history of artificial intelligence, meaning attempts to make computers do things that would seem intelligent if done by a human, goes back to the 1940s. It starts only then because no general-purpose digital computers existed before that time. Virtually as soon as the first such computers were constructed—lumbering, slow, and extremely limited by today's standards—they were used for experiments in artificial intelligence. In fact, the promise of artificial intelligence was a key driver in the

Superhumanly intelligent machines have existed for a long time, decades at the least.

development of computers, both because of the perceived immense usefulness of being able to automate reasoning and learning of various kinds, and because humans have pretty much since the beginning of history imagined that we could re-create intelligence in nonbiological vessels.

Early depictions in fiction of what we might call intelligent machines include Hephaestus's artificial servants in Greek mythology, and the Golem, an intelligent creature made of clay to protect the Jews of Prague in a sixteenth-century story. As mechanical clocks were invented around the fourteenth century, some European clockmakers produced elaborate automatons that could perform complicated movements. These automatons were often installed in churches, where they could impress the townsfolk and visitors. Such automatons had essentially no flexibility, however, and could only perform preset movements. Charles Babbage's analytical engine was a form of sophisticated programmable calculator or early computer, but Babbage himself only produced the plans; actually building the machine was beyond the capabilities of nineteenth-century mechanics.

Digital technology changed all of that. The speed and reliability of digital circuits allowed engineers to construct calculators of unprecedented capability. Initially spurred by the various computational demands of World War II, such as firing artillery at moving airplanes and interpreting radar signals, the first programmable digital

computers emerged in the 1940s and 1950s.¹ While these machines were primitive by today's standards, they could do marvelous things. In particular, they could calculate. At the time, the term *computer* referred to a human (usually a woman) who would perform calculations entirely by hand or with simple manual tools only. Large projects, such as those involved in designing aircraft or artillery for the war effort, could employ hundreds of human computers. The new digital computers could perform these calculations thousands of times faster. This was obviously a revolution in computation, but why is it not seen as the first superhuman AI? After all, calculation is something that we do with our minds, and ever since the 1950s, machines have been better than humans at calculating; the simplest pocket calculator you could buy in a dollar store "thinks" incomparably faster than you do.

Indeed, some people (including some of the inventors) referred to these new computing machines as "electronic brains." But generally, complex calculations in themselves are not thought of as artificial intelligence these days. One reason may be that they are so obviously formalizable: there are clear rules for how to perform a calculation by breaking it down into a series of small and simple constituent calculations. It might be because the process is so obviously mechanizable that we don't see pocket calculators as superhuman AI, even though humans think when they compute. It might also be that the term "artificial

intelligence” was only invented after digital computers, so earlier forms of mechanical “thought” would not count.

But mathematics involves other cognitive tasks than just performing calculations. Mathematicians, after all, do not spend their days doing calculations. Instead, they develop new mathematics. In other words, they prove theorems, and more importantly they reason about which theorems should be proved, using which strategies, to create useful and beautiful new mathematical insights. If we agree that developing new mathematics requires intelligence, a system that can prove theorems as well as a human should qualify as artificially intelligent—and if it does so better than humans do, superhumanly intelligent.

In 1956, Allen Newell and Herbert Simon wondered if they could make a computer prove mathematical theorems. At the time, very few digital computers existed, and they were generally the size of rooms. Newell and Simon chose as their domain proofs in the foundation of mathematics, specifically those in the monumental early nineteenth-century book *Principia Mathematica*, which attempts to ground mathematics in pure logic.² They created a program called the Logic Theorist and tasked it with rediscovering the theorems of the *Principia*.³ In essence, the Logic Theorist is a search algorithm that automates rewriting proofs following well-known rules for equation rewriting, some of which we all learned in high school. It starts from a set of axioms and gradually rewrites them

until it manages to prove what it set out to prove. However, these axioms can be rewritten in an astronomical number of ways, most of which are not interesting. Therefore the Logic Theorist includes some rules or heuristics aimed at mimicking what humans think are interesting ways of theorem rewriting. Using these heuristics, the Logic Theorist managed to rediscover thirty-eight of the first fifty-two theorems in the *Principia*. This achievement was quite remarkable at the time, given the many years it had taken Bertrand Russell and Alfred North Whitehead, both eminent mathematicians, to discover the theorems the first time around.

Given that the Logic Theorist rediscovered these theorems thousands of times faster than Russell and Whitehead and occasionally produced more detailed proofs, it is safe to say that the Logic Theorist is a superhuman AI system. However, the domain in which it is superhuman is very small, essentially only the *Principia Mathematica*. If you want it to prove something in other areas of mathematics, you would have to feed it new axioms, and probably also new heuristics.

The Logic Theorist was developed in the 1950s, so you would expect us to have made some progress since. We have. Sixty-five years of research in the subfield of artificial intelligence called automatic theorem proving have yielded a number of impressive AI systems that have occasionally come up with novel theorems or at least new and

interesting proofs of already discovered theorems. But math is still done by human mathematicians; if you look at a mathematics journal, the papers describing recent advances are written by humans. Chances are, however, that the human mathematicians frequently use modern software such as Mathematica, which builds on many of the advances made in automated theorem proving. They might also use a large language model like GPT-4 somewhere in the process.

It is interesting to note that Herbert Simon, one of the inventors of the Logic Theorist, said that he and Newell had invented a “thinking machine” and had “solved the mind-body problem, explaining how a system composed of matter can have the properties of mind.” This may sound like ridiculous hyperbole in hindsight, given that their program could only do one thing: prove *Principia* theorems. It couldn’t even prove anything about, say, differential equations, much less write a sonnet, draw a picture, or tie its nonexistent shoelaces. But the program was undoubtedly a major advance, and it is easy to understand how the limitations were not easy to see at the time. Isn’t all thinking like theorem proving?

Apart from proving mathematical theorems, what other things do humans do that require intelligence? Playing games, perhaps? Throughout the history of artificial intelligence research, game playing has been an important focus, with some spectacular results.

Chess is perhaps the game that has been most important for AI research. This is not surprising, as this classic board game has long been seen as requiring and exemplifying intelligence in Western culture. Since medieval times, noblemen, students, and clerics have been depicted enjoying the game as an intellectual challenge. In Norse mythology, Odin the all-seeing all-father was described as enjoying a chess-like game. And playing chess well does seem to require sophisticated reasoning. To win against a worthy adversary, you must form a complex plan that is executed in multiple steps, and you must be able to recognize the various threats and opportunities that arise from your adversary's moves. Basically, you need to out-think your adversary. It stands to reason that if someone or something can play chess, it would need to be intelligent. So it is not surprising that early AI researchers were obsessed with creating chess-playing AI.

So how can a computer play chess? The basic algorithm that underlies most chess-playing AI is called Minimax. It basically simulates taking all possible moves for white (if playing white), followed by all possible countermoves for black, followed by all possible counter-countermoves for white, and so on. By assuming that each player will on every turn take the best move, it is then possible to calculate the best possible move to take on any turn. However, calculating the whole tree of possible moves and countermoves for a whole game takes an extreme amount of

time; no existing computer could do it in a lifetime. Actual chess-playing AI systems therefore only use the Minimax search algorithm to a certain depth, perhaps five to ten moves. The search is then combined with a myriad of little tricks, particularly ways of evaluating when a board state is promising.⁴

Since the 1950s, AI researchers kept developing chess-playing AI programs, making steady progress by adding new bells and whistles to the original Minimax algorithm. All along, many confidently predicted that a computer would never play chess as well as the best human players. Why? Well, because playing chess requires intelligence, and computers are not intelligent. And certainly the Minimax algorithm has almost nothing to do with how we understand real biological intelligence to work.⁵

It was therefore a great upset when, in 1997, a computer really did beat the best chess player of the time. IBM invested a prodigious sum of money in building a special-purpose computer and developing the software to go with it, called Deep Blue, and challenged Garry Kasparov, the reigning world champion.⁶ The six games were highly publicized, and the result, where Deep Blue won by one game's margin, sent shock waves through the chess community and beyond. The match was covered prominently in news media across the world, often accompanied by speculation about what the result meant for our understanding of intelligence. Prominent experts who had claimed that

the computer had no chance of winning because it lacked “intuition” or “creativity” were being asked to clarify what they meant, typically having to revise their positions.⁷

However, people generally observed that Deep Blue could do only one thing: play chess. It was no better than the Logic Theorist or a toaster at any of the myriad of other tasks that humans do all day. The experts also often remarked that the Minimax algorithm is nothing like how the brain works. It’s really just a search algorithm. So it would seem Deep Blue was not very intelligent after all. Or at least not generally intelligent. Maybe part of the problem was the choice of chess as a problem domain: given that the game could be played so well by a piece of software that may be sophisticated but at its core was just a search algorithm, maybe chess didn’t actually require intelligence to play?

After Deep Blue, many AI researchers shifted their focus to another classic board game: Go. Much like chess, Go is an ancient two-player board game and occupies a similar place in some Asian cultures as chess does in some European cultures. Unlike chess, games are longer, and many more possible moves exist at every turn. It is also harder to judge how good a board position is: In chess, you can get a reasonable idea of how well you’re doing by counting pieces, but in Go, such piece counting will get you nowhere. Instead you have to analyze the patterns on the board. And indeed, the Minimax algorithm that works so well

on chess plays Go very poorly. Energized by the challenge, many AI researchers inclined to work on games turned to Go as their next big endeavor.

This challenge was overcome in 2016, when DeepMind, a research company that is part of Google, presented its AlphaGo system. Much like Deep Blue, AlphaGo uses a search algorithm at its core. But unlike the Minimax algorithm of Deep Blue, the algorithm in AlphaGo uses statistical calculations to help decide which moves to explore. It also uses several neural networks that assist the search algorithm in evaluating how good a board position is and what move to explore next. These neural networks were trained on logs of expert humans playing Go and then further trained by the system playing against itself.⁸

In a highly publicized event, AlphaGo went on to win against Lee Sedol, one of the best Go players in the world.⁹ A media frenzy ensued and, among the flurry of commentary, quite a few voices warned against the impending AI takeover of everything, as AlphaGo portended the rise of artificial general intelligence. Others noted that Go may not be such a great test of general intelligence after all, and that AlphaGo could really only play Go, just as Deep Blue could only play chess. None of these systems could even play checkers (without reprogramming), much less drive a car or write an essay on witchcraft in medieval Estonia.

Today you can download an app on your phone that plays chess or Go better than any human alive.¹⁰ In other

words, it is safe to say that we have superhuman AI players for Go and chess. Yet people still play these games. In fact, it is now common to train against AI players, and to study the games of top AI systems to learn new strategies and playing styles from them. Within AI research, interest has largely moved on to methods for playing video games (as opposed to board games), and for designing games. I touch on these topics later.

Mathematics and game playing are not the only problem domains AI methods have been applied to. They are also not the only domains where superhuman performance has been achieved. Far from it. To round out the examples in this chapter, let's take a domain of particular practical significance: image recognition.

Being able to look at an image and say what's in it is an obviously useful skill that most adult humans are pretty good at and use all the time. When we're driving or cooking or getting dressed or moving files on our computer, we rely constantly on identifying what we see in front of us. We need to know whether the blue fabric in the drawer is a pair of socks or underwear, and whether the thing moving across the road in front of us is a plastic bag or a cat.

Given the importance of identifying objects (and other forms of image recognition) for everyday living, you might imagine that early AI research would have focused on solving these tasks. Alas, that's not how it went. Perhaps

surprisingly, we had stronger-than-amateur chess-playing and theorem-proving AI before we even had any serious attempts at image recognition systems.

Why was this the case? Partly because early computers did not have enough memory to store images of any reasonable size. But also because we did not have cheap digital cameras that could convert views of the real world into bits and bytes in computer memory. And no computer networks on which to send such digitized images around, so that there would be something to train the AI models on. We also did not have very advanced learning algorithms available, nor did we have processors fast enough to run those algorithms if we had had them. Early computer vision systems instead largely relied on hand-designed detectors and tended to work only under perfect conditions, and not very well even then.¹¹

That all changed in the early 2000s. One crucial component was the ImageNet data set and competition.¹² ImageNet contains millions of images in thousands of labels, and each image is associated with a single label, such as “flamingo,” “harp,” or “volleyball.” Each label contains hundreds of images. These images were crowdsourced, meaning that they were collected and labeled by large numbers of people via the internet. The challenge for the AI methods is to assign the right label to the right image. Starting in 2010, there were annual ImageNet competitions, where researchers sent their best algorithms to be tested against

the data set. At first, the AI systems did poorly, getting perhaps half or three-quarters of their guesses right. But in 2012 a new challenger entered the game: AlexNet, an image recognition system based on deep learning, or neural networks with multiple layers.¹³

A neural network is a machine learning method that is loosely inspired by the structure of the human brain. It consists of a large number of “neurons” that are connected via weighted connections. Usually the neurons are arranged in layers, where the input (e.g., pixels in an image) is presented to the first layer, and the output (e.g., which class of object appears in the image) is extracted from the final layer. Between the first and the final layer there can be a varying number of layers or neurons, from zero to hundreds. Neural networks learn from data by changing the weights between the layers, so that they can learn to reproduce the patterns in the data. Every time the network outputs a wrong prediction (e.g., saying that the image contains a ball when it actually contains a bike), the weights are adjusted ever so slightly in the direction of the correct output. The basic ideas used in neural networks have been around since the 1950s, but it took a few decades to get the important technical details right, and even longer to get the necessary computing power to train them efficiently.¹⁴ Where the neural networks that were studied in the 1950s had dozens of connections, and those studied in the 1980s had thousands, the neural networks that are

trained and deployed today regularly have millions or billions of connections.¹⁵

AlexNet was important not only because it was so much better than the competition—it guessed 85 percent of images correctly—but also because it was one of the first instances of running a deep neural network on a GPU, or “graphics card,” the part of a modern computer that is usually responsible for rendering advanced 3D graphics. The discovery that GPUs made it possible to train bigger and deeper neural networks than before led to something of a gold rush, when researchers from all over the world raced to develop new neural networks that would perform better on ImageNet. And they did. As of the time of writing, the best AI system labels around 99 percent of images correctly, whereas a human who knows the data set well gets around 95 percent correct. (A human who has never seen the images before performs much worse.)

It is thus safe to say that we have AI systems that have superhuman performance on image recognition—at least for particular data sets. A neural network trained on ImageNet might not work as well on pictures from another data set and cannot recognize types of objects that it was not trained on (e.g., cheese slicers).¹⁶ There is some generality, though. In particular, you can start from a network trained on ImageNet and then train it some more (“fine-tune” it) on some other data so that it can do another task, such as recognizing road signs or cheese slicers. But this

network on its own still cannot write poetry or tie its shoelaces any more than Deep Blue or the Logic Theorist can. It also cannot prove theorems or play chess.

Obviously, theorem proving, board game playing, and image recognition are not the only problem domains that have been studied in AI research. They are not even the only ones where we have achieved superhuman performance. There are also many other types of methods than the ones listed here; this chapter is far too short for anything like a representative sketch of the history of AI. To take a few examples, we have not mentioned statistical machine learning, recommender systems, or expert systems. But from these domains and others, it is possible to make some general observations.

The first is the well-known “moving goalposts” phenomenon in artificial intelligence. This refers to how AI research often progresses by aiming to solve some specific problem that people agree would require intelligence to solve, and once the problem is solved, people no longer agree that intelligence was needed to solve the problem. For example, researchers commonly agreed that playing chess well would mean you were intelligent, until Deep Blue. Then the goalposts are moved, in the sense that some other problem is identified that AI methods cannot currently solve. This is not necessarily a bad thing. It’s great to have a set of specific problems to work on to motivate AI progress. And even if solving any particular problem doesn’t “solve intelligence,”

it helps us understand intelligence a bit better. Wouldn't it be strange if we knew which problems were most important to understanding intelligence from the start?

Another general observation is how much of AI progress depends on advances that do not lie strictly within the field of artificial intelligence itself. What made many of the advances described here possible is more data, more processing power, better sensors, and better computer networking. These technologies were mostly developed not for AI but for photography, business communications, video gaming, and other applications.¹⁷ And while we have seen a myriad of advances in AI algorithms, the basic ideas of search, optimization, and machine learning have been there more or less since the inception of the research field.

I have pointed out several times how each of these systems can only do one thing, namely, the very task it was constructed (and trained) for. This point is often lost in the general debate, where the popular media often claim that "AI" (bizarrely used as an uncountable noun, like "water") can do all kinds of things. Not only could the Logic Theorist only prove theorems, but it could not prove theorems outside number theory unless it was reprogrammed. And a neural net trained on ImageNet cannot recognize types of objects that are not in the ImageNet test set unless retrained.

The focus on the term *AI* also obscures the fact that these systems sometimes have little in common. The Logic Theorist and AlexNet, the neural network that

The focus on the term *AI* also obscures the fact that these systems sometimes have little in common.

revolutionized image recognition, are built on completely different principles, and understanding one does not really help you understand the other. They certainly share no program code. We have made some genuine advances, and new AI systems often build on insights provided by older systems. For example, AlphaGo built on both the advances in search algorithms for games that led to Deep Blue and the advances in neural networks that led to AlexNet. But every AI system is built or trained for a particular task (even though this task may have interesting emergent side effects; see chapter 7).

Finally, one thing that all these systems and many others have in common is that they caused journalists to worriedly or excitedly explain that we were close to creating machines that could actually think, that we had found the secret of intelligence, or perhaps even that AI might take over the world. In 2024, it bears pondering that such hopes and concerns were discussed sixty years ago. Are they more relevant now because of all our technical progress? Or is our failure to understand intelligence, and AI systems' failure to take over the world, during the last sixty years reason to believe it won't happen anytime soon?

In the next two chapters, we will look into theories of intelligence to see if we can understand it better, or at least define it. We'll start with theories of intelligence in humans and other animals.

INTELLIGENCE (NATURAL)

What is intelligence? The psychologist Edwin Boring gave a non-boring answer to this question in 1923: "Intelligence is what is measured by intelligence tests."¹ Let's use this quip as a convenient excuse to sidestep (for now!) the hard question of what intelligence actually is, and observe that we regularly measure the intelligence of humans. IQ (intelligence quotient) tests are extremely common; you have probably taken one at some point, perhaps in school, in military aptitude testing, or for joining a club such as Mensa. What are IQ tests, and how do they measure intelligence? Does what they measure correspond to our intuitive idea of intelligence?

Modern IQ tests have their roots in the psychologist Alfred Binet's attempts to measure the mental age of schoolchildren.² Binet sought to identify children who were behind their age group ("mentally retarded") so that

they could receive the most appropriate education. A version of this test, the Stanford-Binet Intelligence Scales, is still in widespread use. In the Stanford-Binet scale, IQ is relative to the subject's age, and the scoring is normalized so that the mean is always 100.³ Many different intelligence tests have been devised, but most of them center on a battery of multiple-choice questions that the test taker must answer. These tests might present a series of geometric figures and ask the subject to identify the next figure, or they might involve reasoning by analogy, grouping words by topic, or finding patterns in numbers.

Intelligence testing is so common partly because it works. Tests such as the Stanford-Binet IQ test have high reliability, meaning that test takers generally score consistently over time. These tests are also predictive of an astonishing number of different things. For example, intelligence testing is pervasive in military recruiting, as test results have shown convincingly that higher-IQ soldiers fight better. They die less and succeed with their missions more often. Even civilians generally live longer and earn more money if they score well on intelligence tests. But some things are more strongly correlated with IQ than others: for example, IQ is not a good predictor of leadership skills.

You may object that IQ is really a rather narrow measure of intelligence. And you would be right. There is much more to life, and intelligence, than completing sequences

of geometric figures. But fret not: psychologists and psychometricians have been hard at work over the last century or so inventing tests of different kinds of intelligence. Moreover, they have studied how these different measures of intelligence correlate with one another.

A leading theory of the structure of human intelligence, which attempts to synthesize these various intelligence measures, is called the Cattell-Horn-Carroll theory.⁴ CHC theory sees intelligence hierarchically: at the top, you have the *g* factor, which can be split into ten broad abilities, which can then further be split into a number of narrow abilities. Together, these abilities constitute human intelligence, or so the theory goes.⁵

Now, what are these broad abilities? Two of them involve memory (*short-term memory* and *long-term storage and retrieval*), and two involve sensory modalities (*visual processing* and *auditory processing*). Then we have *reading and writing* and *quantitative knowledge* (basically, doing things with numbers), and two speed-related abilities: *processing speed* and *decision/reaction time*.

You may argue that these abilities are all a little peripheral to true intelligence. And, indeed, I saved the two most "core" abilities until last: *comprehension-knowledge* and *fluid reasoning*. Comprehension-knowledge is essentially what you know and how well you can communicate and use this knowledge, and fluid reasoning is how well you can reason, form concepts, and solve problems in unfamiliar

situations, what is sometimes known as “thinking on your feet.” These two abilities embody the widely applicable distinction between *fluid intelligence* and *crystallized intelligence*. Crystallized intelligence accumulates as we grow older and wiser (or at least more experienced), whereas fluid intelligence starts diminishing at some point in our adult life.⁶ As a professor in my early forties, I am probably supposed to say something here about how we should not undervalue crystallized intelligence, but my fluid intelligence can’t come up with something clever right now.

At the top of the CHC hierarchy sits *g*, or “general intelligence.” All ten broad abilities and their respective narrow abilities depend to some extent on *g* in the sense that they are positively correlated with it. The correlations here are not incidental; *g* was discovered (or invented) by simply testing a large number of people with a large number of different intelligence tests and using a mathematical technique called factor analysis to find the single variable that correlates best with all of these tests. In other words, the various intelligence tests correlate highly with *g* by construction. This is simply how *g* was defined.

Psychologists mostly agree that *g* is a good construct that represents general intelligence reasonably well. The logic is that because all these different intelligence tests correlate so well with one another (if you score highly on a test of sequence completion, then you are also likely to score highly on a vocabulary test, and so on), they must be

testing something real, and the mathematical construct that correlates best with all the tests should be a good approximation of intelligence. Also, *g* is reliable, meaning that the same person scores about the same when tested at different times. It is also highly predictive of a number of measures of real-life success. For what it's worth, *g* is also around 50 percent hereditary, so the smart thing to do if you want to be generally intelligent is to choose smart parents for yourself.⁷

This being said, there is no shortage of criticisms of *g* and the wider CHC theory. One line of criticism is statistical. If you take any arbitrary set of tests and run a factor analysis on them, you will find the construct that correlates best with them, and chances are that the correlations will be rather high.⁸ So why should we trust this particular set of tests to define what we call "general intelligence" rather than some other set of tests? It is true that the tests correlate pretty well with one another, but maybe they were chosen for exactly that reason. This doesn't mean that they say anything deeper about intelligence. As we will see, the criticism that the intelligence tests are rather arbitrary is hard to shake.

Another line of criticism is that intelligence is less about what you know right now and more about how fast you can learn to adapt to new situations. The tests that are used to define *g* and the various abilities in the CHC theory don't measure how well or fast you learn something new.

As far as I can tell, the reason why tests of learning are not used to define intelligence is that measuring learning ability is slow, expensive, and difficult, whereas an IQ test can be done in less than an hour.

Perhaps even more fundamentally, you may disagree that the best way to define intelligence is by testing how often you answer questions "correctly." You may argue that the type of intelligence that really matters is the ability to come up with new answers or new questions—in other words, creativity. I am certainly sympathetic to this argument myself: great art and music have been created, great theories invented, great stories told, and great companies started by persons who might not have scored well on an IQ test. Many of us have stared at those insufferably boring figure completion sequences and imagined any number of fanciful figures to complete the series, successfully distracting ourselves from identifying the figure that the test constructor deemed to be correct. Alas, it turns out that testing creativity well is really hard, which is probably why it is not well represented in CHC theory. Perhaps we need more creativity applied to the problem of measuring creativity.

It can also be plausibly argued that you cannot have intelligence without emotion. As many have observed, emotion is crucial in human decision-making. David Hume famously called reason the slave of the passions, and modern-day researchers such as Antonio Damasio

Intelligence is less
about what you know
right now and more
about how fast you
can learn to adapt to
new situations.

have compiled ample evidence that emotions affect many facets of our thought processes.⁹ None of the components of *g* measure emotional competence, so CHC theory would seem to be an incomplete account of intelligence if we accept this view.

To complicate matters further, consider another comprehensive theory of intelligence: Howard Gardner's theory of multiple intelligences.¹⁰ Gardner's theory posits nine separate intelligences: musical, visual-spatial, linguistic, logical-mathematical, bodily-kinesthetic, interpersonal, intrapersonal, naturalistic, and existential. As can be seen, this concept of intelligence is considerably broader than that of CHC theory. Unlike that theory, Gardner's theory does not commit to *g* or any other single overarching intelligence. Instead it sees the intelligences as largely separate. For example, bodily-kinesthetic intelligence is not necessarily correlated with logical-mathematical intelligence. The theory of multiple intelligence has had a large impact in some circles and has had several best-selling books written about it. However, it is mostly ignored by mainstream psychologists because of its lack of empirical evidence. Sometimes the theory is brutally rejected: the psychologist George Miller characterizes it as "hunch and opinion," and other psychologists have written similarly scathing indictments.¹¹

While Gardner's theory may have nothing like the rigor or evidence mass of CHC theory, it does raise questions

about the narrowness of *g*. Can a theory of intelligence really be complete if it does not include our ability to manage interpersonal relations or our own bodies? On the other hand, should a theory of intelligence really include musical ability, or is that a little peripheral? The approach of mainstream cognitive psychology is to focus on abilities that can easily be tested and where the tests correlate highly with one another, but you may very well argue that this is a backward way of defining an important concept such as intelligence.

In the last few pages, I have insidiously taken you from a simple answer (intelligence is what IQ tests measure) to what looked like consensus (the *g* factor) to doubt and confusion. In all likelihood, I have left you with more questions than you had when you started reading this chapter. To add to your frustration, or perhaps delight, I am now going to make matters yet more complicated. Let's talk about intelligence in nonhuman animals.

You might think that it is easier to define intelligence in animals because they are, well, less intelligent than humans. Alas, no. There is no standardized test of intelligence for animals, and no agreement on what such a test would consist in. Those clickbait articles about how dogs are smarter than cats (or the other way around) are mostly opinions with only tenuous basis in fact.

The first wave of systematic study of animal behavior was behaviorist psychology. If you walked into a psychology

department a hundred years ago, you would very likely find a bunch of behaviorist researchers doing experiments on animals in their labs (at least if you walked into a psychology department in the English-speaking world). Behaviorists believed that behavior is the response of an organism to stimuli and can be learned through conditioning. In the limit, behaviorism says essentially that conditioning is all there is; we are all born as "blank slates," and everything that seems intelligent in our behavior results from associating actions and stimuli with rewards. The blank-slate idea extends as far as positing that the same learning mechanisms operate in humans and other animals, at least certain "higher" animals. Animals could therefore be used as models of human learning. This was highly practical, as the behaviorist movement heavily emphasized systematic controlled experimentation, including experiments that would be difficult or illegal to perform on humans.¹²

Many of the most famous psychology experiments from the behaviorist era were animal experiments, including Pavlov's experiments on classic conditioning, where he taught dogs to associate the ringing of a bell (the stimulus) with food (the reward).¹³ Famous behaviorist animal experiments also include numerous experiments on teaching rats to navigate mazes, and B. F. Skinner's work on comparing reward schemes where pigeons are taught to pick at buttons (actions) and receive food pellets as rewards.¹⁴

Some of this work has provided scientific insights that are still valid today, though the behaviorists' perspective on intelligence seems oddly narrow to a modern reader.

One of the many reasons that behaviorism unraveled is that the assumption that these mechanisms of learning are universal did not hold up. For example, you can easily teach a rat to associate a sound with an electric shock, but you cannot teach it to associate a smell with the same shock.¹⁵ From a behaviorist perspective, this makes no sense, because both the smell and the sound should be stimuli. But it turns out that rats, by virtue of their evolutionary history, have a much easier time learning some things than others. (Presumably being able to associate sounds with pain was more useful for the rats' ancestors than being able to associate smell with pain.) This is true for other animals as well. They are wired to learn some things more easily than others and in fact come with a whole host of prewired behaviors. This calls for studying animal behavior in a way that takes the specifics of each species into account.

These days, animal behavior is mostly studied not as a subfield of psychology but as its own field, ethology. Ethology relies heavily on insights and perspectives from evolutionary biology. Behavior is seen as an adaptation to the specific ecological niche of a species. A lion is a group-hunting apex predator, a whale shark is a mostly solitary filter feeder, and an ant is a member of a highly collective

species with highly delineated roles among individuals. For all these species, their behavior results from a combination of genetically determined development, individual learning, and sometimes cultural transmission.

Because each animal species has evolved the capabilities it needs to survive and reproduce in its particular ecological niche, the sets of capabilities differ widely between species. From the perspective of human cognition, many animals would therefore seem to have extremely unbalanced and nongeneral capabilities. As the ethologist Frans de Waal puts it, "Ranking cognition on a single dimension is a pointless exercise. Cognitive evolution is marked by many peaks of specialization."¹⁶

It is striking that some animals have cognitive abilities that rival or even surpass those of humans in particular cases. Sometimes other animals outperform humans even on what we usually call "core" cognitive abilities, such as memory and strategic reasoning. For example, chimpanzees have been shown to outperform humans in abstract choice games, getting closer to the game-theoretic optimum.¹⁷ Similarly, pigeons do better than humans at the Monty Hall dilemma, a classic puzzle.¹⁸ Corvids, a family of birds that include crows and ravens, have astounding spatial memories because they need to cache foods in preparation for winter. For example, a corvid species called Clark's nutcracker regularly caches food in 20,000 to 30,000 different locations and remembers around 90

percent of these locations after a year.¹⁹ I'm envious, as I frequently don't know where to find some ingredients when cooking, not to mention the TV remote.

Cognition is intrinsically connected to perception, and many species of animals have very different sensory capabilities than we have. Humans are a predominantly visual species, with good hearing in certain frequencies and a limited sense of smell.²⁰ We don't need to look further than to our most common household pets, cats and dogs, to find a different battery of sensors and sensory-based cognition. When a dog finds a person in a crowd or in a vast forest after only smelling a shirt that person has worn, this is not only a sensory feat we can't get anywhere near but also arguably a cognitive feat we can't replicate. We don't even know what it's like to trace paths based on smells that may vary in intensity and be carried by winds and obscured by other smells. Some animals even have senses that we completely lack, such as the electrical-field sensors of many fish and the echolocation of bats. What kind of intelligence do you need to use echolocation to capture flies in midflight?

If we look at animals that are evolutionarily or ecologically even further removed from us, the "unevenness" of cognition stands out even more starkly. Octopuses are animals that are so different from us in every respect as to practically be aliens, but they have large brains and display examples of stunningly intelligent behavior. Much of this

behavior is connected to manipulation of objects and their own bodies. Octopuses can unscrew jar lids, disassemble crabs, squeeze through tiny openings, and devise complicated schemes to escape their confinement tanks. Some of them are so good at this as to be almost uncontainable. Certain species of octopus can change the color and shape of their bodies almost instantly to camouflage themselves against a wide variety of backgrounds. On the other hand, most octopus species seem to have no social capabilities at all, with no observed nontrivial communication or learning between individuals.²¹

In a completely different corner of the multidimensional space of cognitive abilities are ants. Individual ants are seemingly too small to have room for much of a brain, and in practice their behavior seems rigid and not overly complex. Ant societies, on the other hand, are highly complex, characterized by advanced division of labor and the ability to coordinate on building impressive structures. Anthills can be hundreds of years old and contain hundreds of thousands of ants. Through a sophisticated communication system based on pheromones, these individually rather simple creatures manage to build structures vastly larger than themselves and, even more impressively, manage the logistics of keeping all the ants and larvae fed and cared for. You could argue plausibly that the actual intelligent organism here is not the individual ant but the ant society. Of course, you could make the same argument

Octopuses are animals that are so different from us in every respect as to practically be aliens, but they have large brains and display examples of stunningly intelligent behavior.

about humans, as our companies, states, and other organizations are capable of doing things no single human could.

Seeing the very different capabilities of these different species, there does not seem to be any reasonable way we could measure their intelligences on the same scale. Not only is there not a practical test we can subject the various species to (and that they would agree to take), but it is also not clear how we would rank them even if we could test their intelligences. Which is more intelligent: a nutcracker, an octopus, or a mouse? Each has capabilities that the other species don't have. Even if we limit our comparisons to primates, their cognitive abilities are largely incommensurable. Our closest relatives, chimpanzees and bonobos, are physically similar but entirely different in social structure.

What does all of this tell us about human intelligence, and the possibility of artificial general intelligence? One response is that while animals have highly varying capabilities, and it is hard to say that one species is more intelligent than another, humans play in a different league, because we have general intelligence. To see why this could be true, consider the following thought experiment. Take an octopus or a mouse or a corvid and increase their intelligence by a hundred times, in the same way we make a computer "more intelligent": make it think faster and have more memory. So each of these animals would be the same as before, except it thinks a hundred times faster and can

remember a hundred times more. Most likely they would just become better at things they are already good at. A super-octopus would have amazing dexterity but would probably still have no social skills to speak of; a super-corvid would be able to remember a staggering number of things and places but would probably still not learn to read and write. These super-animals would arguably still lack many aspects of human intelligence.

However, this human exceptionalism is at odds with everything we know about how evolution works. Evolution tends to be gradual and to build scappily on what is already there. New anatomical or cognitive features are usually built on top of or repurposed from existing organs or mechanisms.²² We know that the cognitive abilities of our ancestors were developed gradually over many millions of years.²³ Why would a completely new, "general" intelligence suddenly evolve? It is much more likely that our intelligence is a collection of special-purpose capabilities that have evolved in response to specific needs, just as it is for all other animals we know of.

We can also turn the previous intelligence argument on its head. Imagine that we took a human (such as you or me) and made a super-human who thought a hundred times faster and could remember a hundred times more. This super-human would still probably not be able to track a person by scent like a dog, and definitely not be able to echolocate like a bat or blend in with the environment by

changing skin color like an octopus. Yes, these capabilities depend on anatomy humans don't have, but they also depend on cognition. It is not clear that the super-human would remember 20,000 seed caches in a forest, because the super-human might not be very interested in birdseed or forests, and humans are bad at remembering things they don't take an interest in. Furthermore, even this super-human would arguably lack whatever other cognitive capabilities that various animal species may have but we don't know about because we have never looked for them, because we don't have them. We can call them the unknown unknown forms of intelligence. (There may also be forms of intelligence that neither we nor any other animals have but which an organism could conceivably have.)

Some readers might at this point experience a certain frustration. It would seem obvious that humans are the most intelligent species on the planet. After all, humans went to the moon and invented computers. Crows did not. Shakespeare and Bach were (supposedly) human and not chimpanzee. This is indisputable but assumes a human perspective on which feats we consider important. Space-ships, sonnets, and sonatas may not impress an octopus—why would they care? We know for sure that cats don't.

In sum, the human study of human intelligence has come a long way, and we have sophisticated theories with strong explanatory power in some domains. However, these theories of intelligence are rather narrow and

exclude much of what we consider important as humans. More inclusive theories exist, but they are decidedly less rigorous. From the perspective of animal cognition, human intelligence would seem like yet another idiosyncratic set of cognitive abilities that were evolved to allow us to survive and reproduce in a particular ecological niche. This idiosyncratic set of cognitive abilities allowed us to build complex societies that maximize our individual agency, making us feel like kings of the world. But that in itself does not necessarily mean our intelligence is very general.

INTELLIGENCE (ARTIFICIAL)

In the previous chapter, we examined thoughts about the nature of intelligence, and particularly general intelligence, in humans and other animals. A good deal of thinking about intelligence and general intelligence has also come from within the AI community, which we will explore in this chapter.

If there is one idea about intelligence from AI research that has captured the public's imagination, it is probably the Turing test. Alan Turing, war hero, mathematician, and computer science pioneer, proposed what he called the "conversation game" in his 1950 article "Computing Machinery and Intelligence."¹ The Turing test is conceptually simple: a human judge communicates with both a human and a computer and needs to figure out which is which. The computer "wins" if the judge cannot tell human from machine. In Turing's paper, the communication takes

the form of typing on keyboards and getting answers via a printer. Today we can just think of this as texting or chatting. (In fact, Turing's whole article is amazingly prescient, full of interesting ideas, and well worth reading today.)

The Turing test has become extremely influential but is also subject to much criticism. One problem is that it depends on the skills of the judge, and not only are some judges better than others but ordinary people are getting better at judging Turing tests all the time. In 1964, Joseph Weizenbaum created a program called Eliza as a computational satire of Rogerian psychotherapy, a school of therapy that, according to detractors, mostly just rephrased whatever the patient said as a question.² Eliza is what we would now call a chatbot, letting you type lines of text and answering with text. The program is very simple. Mostly, it scans the text you type for certain trigger words and uses them to ask you questions. When Eliza cannot find the right words or structures in what you write, it will give you generic replies such as "Can you elaborate on that?"; but when it detects that you mention, for example, your mother, it will ask you to tell it more about your mother. Yes, it really is that simple.

The interesting part from a Turing test perspective is that, at the time, quite a few people who tested Eliza thought it was actually a human. Even when Weizenbaum explained that it was just a computer program, some were incredulous. They could not imagine a computer showing

such understanding of humans and being able to generate such coherent text. Some even suggested that Eliza could be used for actual psychotherapy, much to Weizenbaum's consternation. Today, Eliza would not fool a high school student, maybe not even a middle school student. Everyone who is at least marginally online will have interacted with plenty of chatbots and other text-producing computer programs and will have learned to spot them. Most people who spend some time with any conversational AI system quickly learn the chatbot's "style" and failure modes and know what to say to differentiate a human from a computer program. This goes for the fancy new ones such as ChatGPT as well. Once you have played with it awhile, you learn to recognize its tricks and mannerisms. Because the Turing test depends on a human judge, and human judges vary and learn quickly, this seems like a serious issue. How could a test or definition of intelligence be relative to an observer?

Perhaps more fundamental issue with the Turing test is that it is quite narrow: it only really tests conversational fluency. Many aspects of intelligence are not used when carrying on a conversation, especially a conversation via text. We know that existing AI systems can perform super-human feats in specific domains, such as chess or image labeling, without being generally intelligent. Maybe it's the same for conversations? After all, being good at conversations may just be another specific domain. The emergence

of large language models with impressive conversational abilities but few other capabilities might suggest that this is the case. We humans have built much of our culture and society around written communication, but it is possible to live a full life, and to solve plenty of cognitively challenging tasks, without writing. Surely humans who cannot speak or type can still be intelligent. So it would seem that a test or definition of general intelligence in machines would need to take performance on a wide range of tasks into account.

So here is an intuitive definition and test of (artificial) general intelligence: the general intelligence of someone or something (an *agent*, in technical language) is its capacity to solve all possible problems. Everything, from composing sonnets to sorting lists to tying shoelaces to balancing the state budget of Italy, and so on. Of course, any given agent is not likely to do well on most problems; for example, I know nothing about cross-stitch and very little about accounting, and you may be entirely worthless at identifying edible fish in the Amazon. As am I. It would be impossible to actually test an agent's general intelligence because the number of possible problems is infinite. But intuitively, the idea that general intelligence is your average ability to do *everything* makes sense.

You can even define this formally and write a computer program to test general intelligence in this way. We are not going to dive into the formalism here—this is a book for a

general audience—but let me sketch out the concepts for you. Every problem is an *environment*, and an agent takes actions in that environment. For every action the agent takes, it receives an observation of the new state of the environment, and a reward. The reward could be a positive or negative number or zero. This is called the *reinforcement learning* framework and is adapted from behaviorist psychology, which was in vogue in the early twentieth century, as mentioned in the previous chapter.³

An example of an agent acting in a reinforcement learning framework is a pigeon learning to pick at a certain button, but only when a green light is lit. Every time the pigeon picks the button, it gets a delicious seed if the green light is on (positive reward); but if the green light is not lit, then the pigeon gets a punishment, such as an unpleasant sound. (These kinds of experiments were at the core of behaviorist psychology.) As another example, shoelace tying could be an environment. The state of the environment at any given time involves the position of the agent's fingers and the shoelaces, and the agent gets a reward if or when it manages to tie the laces, with a higher reward if the knot can be untied later but will not untie itself. Yet another example would be a neural network-based agent playing a video game, where the score of the game might be used as a reinforcement signal.

To take a more complex example, we can think of you doing your job in reinforcement learning terms; if you do

your job well, you might get a fat bonus at the end of the year, whereas if you mess it up royally, you might have to look for other employment opportunities soon. (The extent to which employers engage in this kind of reinforcement varies widely, with financial services at one end of the spectrum and public sector bureaucracy at the other.) One of the differences between you doing your job and a pigeon picking at buttons is that you typically get your reinforcement with much more delay. This is consistent with your (presumably) being a more sophisticated agent. In particular, you can form intermediate representations of the effects of your actions from your observations of their effect on the environment and make your own inferences about how they may contribute to your future reinforcement. If you take today off, your next delivery might be late, lowering your expected bonus.

Another concept is the *return* of an environment, which is simply the sum of all the reinforcement an agent gets from that environment, positive as well as negative. If the agent solves that environment well, it will get a high return.

So how does this help us understand general intelligence? Like this: imagine subjecting an agent to *all* possible environments—everything from video games to picking at buttons to herding cats to doing your job in finance or public-sector bureaucracy or whatever it is you do. Because some environments can go on forever, we can

set a time limit and only count the return the agent gets within that time in that environment. We can then define the general intelligence of an agent as its average performance on all possible environments.

While it would obviously be practically impossible to test anyone's or anything's general intelligence in this way, as there are an infinite number of possible environments, the definition makes intuitive sense. What is generality if not the ability to handle diverse environments? And defining intelligence against all possible environments is simply the most general form of generality.

However, there is a problem with this idea. This problem might at first seem like a trivial technicality, but it is actually fundamental. The problem is the "no free lunch" theorem. As you might guess, its name alludes to the saying "There's no such thing as a free lunch." What the theorem says is that "any two algorithms are equivalent when their performance is averaged across all possible problems."⁴ In other words, *all* agents would have *the same* general intelligence—even the really stupid ones that just do random things.

This is highly counterintuitive, so let's unpack it. For every problem, you can imagine the opposite problem, where the best strategy to solve the first problem is the absolute worst strategy to solve the second problem—so bad that even taking random actions would be better. It turns out that when you formalize the notion of a problem, it can be shown that for every problem, there is a corresponding

What is generality if not the ability to handle diverse environments? And defining intelligence against all possible environments is simply the most general form of generality.

opposite problem. This means that the good performance of an agent on a problem is canceled out by the bad performance of the agent on the opposite problem. This has been shown to be true for optimization, search, and supervised learning problems.⁵ The most important assumption is that the algorithm has the same amount of time to solve each problem.

How can this be true, when we know that some algorithms are in practice better than others? Because in practice, we never test an algorithm on all problems. We don't even test algorithms on a representative sample of all possible problems. We test them on the kind of problems we encounter in the real world or find important or interesting for some reason.

The no free lunch theorem has the consequence that if we want to define general intelligence as performance over a very large set of problems/environments, we need to find a way to limit this set so that it is smaller than all possible environments. Alternatively, we could make some problems/environments count more than others. If we want our definition of general intelligence to be relevant to our concerns as humans, the choice of problems/environments should reflect this.

One interesting attempt at circumventing the no free lunch theorem in this way is the concept of universal intelligence by Shane Legg and Marcus Hutter.⁶ Universal intelligence is defined as the average reward over all

possible environments, but weighted so that simpler environments count more. “Simpler” here doesn’t necessarily mean “easier.” Instead it refers to environments that are easy to describe. These environments also favor a certain kind of intelligent behavior, because Occam’s razor (to favor the simplest hypothesis that explains a phenomenon) tends to be a good assumption for an agent. An appealing aspect of universal intelligence is that it can be formally defined, using the minimum length of the description of an environment as a measure of the environment’s complexity. A less appealing aspect of this formalization is that it is provably impossible to know the length of the shortest description of a program, and therefore of an environment. Also, universal intelligence still requires testing on all possible environments. So this is very much a theoretical, rather than a practical, definition.

Having a mathematical definition of intelligence makes it possible to prove various things about it. For universal intelligence in particular, we have a model called AIXI, which in theory implements the optimal agent, that is, the agent with the highest possible universal intelligence.⁷ Just as universal intelligence is unmeasurable in practice, so AIXI is not a practically implementable algorithm. But it serves as inspiration for various attempts at building useful AGI algorithms.

While universal intelligence has nice theoretical properties, we can ask whether it really measures what we want

to measure. Do humans score highly on universal intelligence? This is not clear at all without getting a sense of what the space of all possible environments, or even all possible "simple" environments, looks like. It is possible that most of these environments would look quite alien to the typical human, and even to atypical ones such as you and me. If humans turn out to have low universal intelligence, what would that mean? Would it mean that we have the wrong definition, and general intelligence is something else? Or would it mean that humans are not generally intelligent?

These questions are hard to answer partly because it is difficult to visualize what an infinite space of possible problems or environments would look like. You may have made your way through the past few pages thinking that this is all terribly abstract, and you would be right. But taking the word *general* in artificial general intelligence seriously would seem to require a definition that somehow touches on all possible environments. On the other hand, one may wonder why humans, evolved to live in a particular ecological niche, should be able to perform well in all possible environments, as those are not part of that niche. From the standpoint of evolutionary biology, it is very hard to explain why we would have evolved to perform well in environments we did not evolve in.

An underlying assumption of universal intelligence and similar theories is that what agents do is solve prob-

lems or perform in environments that provide a reward, which is a positive or negative number. The argument for this assumption is simply that it is convenient. There is a perspective from which we can see all goal-directed activities as performing in environments that provide rewards. However, it is often quite a stretch. When you think about it, relatively few of our activities provide unambiguous rewards within a limited time frame. How do we quantify the reward for making a nice dance move, finding a thoughtful gift for your friend, or having a nice walk in the afternoon? When does the task of sitting well in your office chair start and end? If you perform this task badly, you may or may not be rewarded with a bad back at some point. A big part of our reward for eating healthy and driving well is that we might live a few years longer; but this reward comes only at the end of our lives, so the time needed to complete the task is all our lives. And even then, we run into what machine learning refers to as the credit assignment problem: say that we are “rewarded” by living to ninety-five years of age instead of ninety-one. What was that a reward for: eating habits, driving habits, good performance at work, good skills at napping? The answer is almost certainly some unknowable combination.

We can follow this reasoning all the way to the end and simply equate an agent’s general intelligence with its evolutionary fitness. Roughly speaking, we can approximate an individual’s evolutionary fitness as the number of

grandchildren that individual has. But this definition has a number of strange consequences, such as that bacteria are the most generally intelligent beings, because way more bacteria than animals exist on Earth, as measured both by number of “agents” and by weight.⁸ Clearly, evolutionary fitness on its own is not a good measure of general intelligence.

Researchers have made various other attempts at defining general intelligence, seeking to go beyond the subjectiveness and shallowness of the Turing test and the abstraction and impracticality of universal intelligence. One noteworthy example is François Chollet’s work. Chollet defines intelligence in a way reminiscent of universal intelligence, but with several caveats. The informal version of his definition is as follows:

The intelligence of a system is a measure of its skill-acquisition efficiency over a scope of tasks, with respect to priors, experience, and generalization difficulty.⁹

The first interesting thing to note about this definition is that it centers on the ability to learn to solve tasks rather than on the solving of tasks itself. The other interesting thing is that it takes into account priors and experience, meaning that if you are a human being who earned a PhD, your skill-acquisition performance will be judged

differently than if you are a human being who did not go to school, or for that matter if you are an octopus or an artificial neural network. In the terms of CHC theory, it focuses on *fluid intelligence* rather than crystallized intelligence.

Although this definition is still abstract, Chollet also proposes a more practical benchmark for general intelligence building on the definition. The Abstraction and Reasoning Corpus (ARC) consists of a large set of picture sequences, superficially similar to a standard IQ test. However, the rules governing the picture sequences are significantly more complex and also more different from one another. ARC is thus geared to testing fluid intelligence; each sequence is essentially its own problem. However, it is not clear whether the kind of fluid intelligence required to perform well on ARC is the same as that needed for other tasks.

Another interesting attempt at formalizing intelligence is the “work on command” framework.¹⁰ This framework sees general intelligence as the ability to fulfill a large number of requests rather than as performing in a number of environments; the tasks are specified in a language understandable to humans, which biases the distribution of tasks to be more related to human interests and capabilities.

As we have seen in this chapter, AI researchers have tried to define (general) intelligence in a number of different ways. While some definitions overlap, major differences

It seems very possible to create a definition of general intelligence that humans would not score very highly on.

also exist. The more abstract definitions have a certain appeal, but the no free lunch theorem points to the impossibility of completely general intelligence. Given this startling result, it is unclear whether these more abstract definitions focus on the range of skills and tasks that are relevant to us humans. It seems very possible to create a definition of general intelligence that humans would not score very highly on. On the other hand, the more concrete tests, like the Turing test and ARC, seem rather narrow in relation to the vast breadth of capabilities that humans have. Finding the right balance seems tricky. Maybe it is a matter of what you want to use such a definition for, or in other words, what kind of artificial general intelligence you are looking for.

VARIETIES OF ARTIFICIAL GENERAL INTELLIGENCE

It is time to return to the title subject of this book, artificial general intelligence (AGI), and try to explore what it might mean. AGI is used by a wide range of people to signify some kind of development of artificial intelligence to come at some point in the future. These AI systems would have more general abilities than those that exist today. Beyond this, people who use this term seem to mean quite different things. Some have clearly thought more about what they mean when they say AGI than others. Some who use the term have little technical knowledge; others are AI researchers themselves. Some who use the term have little knowledge of how human societies work; others are economists, sociologists, politicians, or business leaders.

To understand the many different things people mean when they talk about AGI, let us try to draw out the ways in which different concepts of AGI differ. These can be seen

as dimensions along which concepts of AGI can vary. We could use these dimensions to organize and compare different visions of AGI. Because no actual AGI exists, we cannot use examples from the real world to illustrate these ideas, so we will use the next best thing: examples from science fiction.

Range of tasks: An AGI system should be general, and this is often understood as being able to complete a wide range of tasks. But exactly how wide a range? We know that being able to perform well on every logically possible environment in a fixed time is impossible, per the no free lunch theorems. But that is arguably not so relevant, as most logically possible environments are uninteresting to us anyway. So which tasks should an artificial general intelligence limit itself to? Is being able to solve a hundred reasonably separate tasks sufficient to be called general? How about a thousand, or ten thousand? And how different do these tasks need to be? The droids in the Star Wars universe, such as R2-D2, seem to be very competent, but only in limited domains, such as hacking and repairing computer systems. Would that be general enough to be called AGI?

Domain of tasks: Can an AGI system be limited to a particular domain and still be called a general intelligence? For example, if we had a system that could play almost any video game in existence and quickly learned to play games it had never seen, but could do nothing else, would

that still qualify as an AGI system? If it could survive as a hunter-gatherer in the Kalahari Desert but do nothing else, would it be AGI?

Embodiment: Would it be possible for an AGI system to be disembodied, meaning that it has no particular sensors or actuators of its own, and no body to get around with? Science fiction offers many examples of AGI systems tied to a specific body, either a somewhat humanlike body (like the Terminator or WALL-E) or a very different body, such as the spaceship that HAL controls in *2001: A Space Odyssey*. But there are also examples of AGI systems that exist only virtually, distributed over a great many networked machines. The majority of the most impressive AI systems today are entirely disembodied, running on interchangeable computers in large faraway data centers, and only generate text or images in response to other text or images. Real-world robotics is proving to be much harder than operating in fully virtual domains, yet some people argue that an AI system needs to be able to function in the real world to understand what it does.

Humanlikeness: How similar in capabilities does a system need to be to a typical human to count as AGI? Does it need to be able to do everything a human can do? Does it need to perform these tasks in a humanlike manner? It is entirely conceivable that an AI system could perform a large variety of tasks, including highly complex tasks, but have minimal overlap with the tasks that humans can

perform. Would this count as AGI? Many people's implicit definitions of AGI seem to assume that the system is very humanlike, although perhaps better. Science fiction has plenty of examples of robots that are so humanlike in both appearance and capabilities that they regularly fool humans, such as Demerzel in the *Foundation* TV show or the various robots in *Westworld*. Good examples of intelligences whose capabilities are radically different from ours are harder to find, as these are by necessity harder to conceptualize. Several of Stanisław Lem's stories, including *Solaris* and *Fiasco*, attempt to describe minds that are extremely different from ours.

Zero-shot or training: Does the AI system need to be able to perform all the tasks right out of the box? Or do we mean that the system should be able to *learn* to solve any task? How much time would the system get to learn to do new tasks? For this question, the extreme answer of "no learning needed" is implausible, as many tasks need to be learned through trial and error. Think of, for example, playing a new video game where you don't know the rules; any intelligence would need to fail a few times when learning to play it. The AIXI algorithm described earlier is an example of the opposite extreme, where the agent has an infinitely long time to learn. But an agent that might take the lifetime of the universe to learn a new skill is most likely not what people talk about when they talk about AGI. One reasonable middle path could be that the AI

Good examples of intelligences whose capabilities are radically different from ours are harder to find, as these are by necessity harder to conceptualize.

system would need, at most, as much time as a smart human to learn every task.

One agent or many: It is common to see AGI referred to as a single, unified agent, much in the way we think of a human. This agent would have a single consciousness and a single set of values; it knows itself and what it wants. We often see depictions of such AGI systems in science fiction, such as HAL in *2001: A Space Odyssey* or the Terminator. Many of the arguments in Nick Bostrom's book *Superintelligence* implicitly assume that the AI system is a single, coherent agent.¹ Alternatively, as we have seen, general intelligence can be thought of as the property of a collective rather than an individual—perhaps even of a society. Social insects such as ants are one recognizable natural paradigm for collective intelligence. Could an AGI system be built out of a large set of independent agents that are themselves less than artificially generally intelligent?

Intentional agent or services: When thinking about what an artificial general intelligence would be like, one naturally imagines an agent that has its own intentions and acts on them out of its own free will. But when one then looks at the various software systems that exist today, one finds a glaring lack of intention and independent action. Even the most sophisticated AI-powered systems are merely tools or services; they are built for a particular purpose and deliver some kind of output in response to human input. The question in the context of AGI is whether

a system that is only a collection of services counts as AGI, or whether it needs to have its own intentions and act independently on them.

By asking these questions, we can create a taxonomy of meanings of the term *AGI*. While we could enumerate all possible combinations of the dimensions of variance, this would be a very long list where most combinations would not be very interesting. Let us instead look at some of the common conceptions, with a few remarks on each. We will continue to use examples from science fiction to exemplify these concepts where appropriate.

Human-level AI: This concept of AGI is a machine that can perform more or less the same range of tasks as a human at more or less the same level; it might be much better than humans at some of them. For examples in science fiction, think of the character Bishop in the movie *Aliens*, Data in *Star Trek: The Next Generation*, or Ava in *Ex Machina*. While a human-level AI would not necessarily need to look like a human, it would seemingly need to be embodied given how much of human intelligence is grounded in our bodies. A human-level AI would presumably also learn to perform tasks on the same timescale as a human. The human-level AI concept of AGI is easy to imagine because it is “just like us, but AI.” But humans are evolved to fill a specific ecological niche and have an idiosyncratic set of capabilities that fill this niche. It’s not clear why we would choose to build a system with exactly these capabilities. It

is also probably much harder to build a system that has the exact same range of capabilities as a human than a system that is wildly superhuman at some tasks but unable to do many other tasks that humans can do.

Special-purpose droid: Another common version of AGI in science fiction is a special-purpose intelligence that is significantly better than humans in some particular domain but generally worse or lacking in other respects. These agents would have intentionality and some kind of generic reasoning skills, giving them some capability for independent action and a significantly wider skill set than any AI system that exists today. At the same time, they would not be able to handle many situations that humans often face. Good examples in fiction are the droids in the Star Wars franchise: R2-D2 and C-3PO have apparently temporally extended self-consciousness, intention, personality, and reasoning skills within their specific niches but are generally inept in most situations a human would face. (Perhaps this is how certain other intelligences would see humans.)

Disembodied mind: Iain M. Banks's Culture novels are set in a utopian civilization where humanlike beings coexist with *Minds*, enormously intelligent machines. Minds don't have bodies of their own but are largely responsible for keeping things running in society and thus control a large variety of mobile robots. Banks envisions Minds as having many humanlike traits, including empathy and a

sense of humor; they have a great deal of intentionality. Basically, they are much like humans, only with a thousand or a million times greater memory, attention span, precision, and processing speed. A Mind has an enormous knowledge bank but must still acquire knowledge as we do, through communication and observation. A Mind is not omniscient. The novels in the Culture universe feature many examples of Minds not knowing what to do or how to do it, because they don't have the requisite knowledge.

Omniscient being: What if we say that the range of tasks is all tasks, zero-shot (no training needed)? This means that the AGI system would instantly know how to perform any task. Arguably, this would mean that we had created an omniscient being; while not omnipotent in a physical sense, it would be omnipotent in a cognitive sense. Luckily (?), we know that such omnipotence is not possible (within fixed time spans) per the no free lunch theorem. This kind of AGI belongs in theology rather than in scientific discourse.

Alien intelligence: What would it mean to completely remove the requirement of humanlikeness? Let's say that an AGI system could perform a wide range of tasks, much wider than what any AI system today can perform, but those tasks have almost no overlap with what a human can do. This would be a system with a completely different set of capabilities from ours—maybe as different from our capabilities as those of a honeybee from a mantis shrimp.

Effectively, such an intelligence would live in a different world. Its goals and intentions would, to the extent it had intentions in the way we conceive them, be largely unknowable to us. It might also be useless to us because it might not be able to do the tasks we need help with, or understand what we need, or simply care. Imagining what such a system would be like is almost by definition extremely hard. The closest we can get is probably the many first-contact stories written by various science fiction authors. As mentioned earlier, the experience of first contact with a truly alien intelligence is a central theme of Stanisław Lem's work. It is often not clear in Lem's stories whether we experience a biological or machine-based intelligence, but the being's thinking is in some sense orthogonal to ours. China Miéville's stories also feature examples of very alien and fundamentally unfathomable intelligences, for example, in *Perdido Street Station* or *Embassytown*. One might argue that to the extent an AI system is built by humans, it will not be truly alien, but as we will see in chapter 8, an open-ended learning system might learn from a self-created world that is substantially different from the one we inhabit and therefore learn skills that are very different from ours.

Collective intelligence: We don't have to require that the AGI system be a single agent, with a unified belief system and access to the various parts of its own processing. It might be a set of smaller systems that are logically or

perhaps even physically separate but able to communicate in some way. The individual parts or subsystems might be similar to one another or very different and have different areas of responsibility. The decisions made by this kind of AGI system would likely reside not in a particular subsystem but rather in the system as a whole, much as no individual ant decides the position and shape of an anthill. An example from science fiction is the Borg in *Star Trek: The Next Generation*.

Clippy on steroids: What could AGI look like if we jettisoned not only the embodiment and humanlikeness requirements but also the requirement that it have intentionality? One answer is a collection of “intelligence services,” much like today’s productivity software. Imagine Microsoft Office, Adobe Creative Suite, or Google Workspace, but with the capability to do basically anything that a suitably trained human would do (and which could be done remotely). It would give you suggestions for what you might want to do (Write a story? Go on a trip? Do your taxes? Rethink your life choices?), and assist you with every part of it or do all of it for you, if you prefer. Communicating with such an AGI system could be done in fully natural language, and other interfaces as suitable; perhaps new interfaces could be invented on the fly. For all its capability, this system would not have any will of its own and would simply act as an almost infinitely powerful set of tools for its users to work with, play with, or whatever they

may want to do. An example in science fiction that comes close to this is the ship's computer in *Star Trek: The Next Generation*, including its Holodeck functionality.

The example conceptions of AGI that I have listed are by necessity just a sample; reading through the last few pages, chances are you have thought of some other variation on the AGI theme that I did not list. For example, many possible intermediate positions exist between the special-purpose droid and the disembodied mind. As you can see, these visions of AGI are starkly different, and in some cases even incompatible with each other. Yet, in the general discourse, the term *AGI* is often used without much qualification, meaning that we don't know which particular vision it refers to. This obviously affects the estimates of timelines for AGI or statements about whether it is at all possible. Saying that we would have something like a special-purpose droid in twenty years is very different from saying that we would have a disembodied mind in the same timescale. And while many would argue that we will never achieve the omniscient-being variety of AGI, many of the same people would agree that the Clippy on steroids vision of AGI is achievable in some timescale.

More crucially, commenters who use different concepts of AGI may disagree not only on timelines and feasibility but also on the direction of AI development toward AGI. The kind of effort that would lead to Clippy on steroids might not help us move toward human-level AI, and

Commenters who use different concepts of AGI may disagree not only on timelines and feasibility but also on the direction of AI development toward AGI.

alien intelligence and collective intelligence might be different directions as well.

This book is not committed to any one concept of AGI. In other words, I am not pinning down what AGI “really” means, because it means different things to different people, and no consensus exists. I will keep talking about AGI systems as systems that are significantly more generally capable than any of the systems we call artificial intelligence today.

One thing I have not discussed in this chapter is what underlying principle or architecture an AGI system uses. Some people would choose to define AGI not in terms of a system’s capabilities but in terms of its manner of operations. In particular, some researchers argue that AGI can in principle not be built only on the type of inductive learning that is performed by deep neural networks.² A related question is whether AGI will be clearly distinguishable from the approaches to AI we have right now, or whether something completely new will be required. I will leave these questions open as we explore practical roads to AGI.

Having tried to give an overview of the various ways that the term *AGI* is understood, and having failed to find a universally agreed-on concept, I think now is the time to get technical. The next three chapters consider technical approaches toward AGI. Chapter 6 reviews AI approaches to creating more general systems, whereas chapters 7 and 8 describe two particular approaches to more general AI in some more detail.

PRACTICAL AGI DEVELOPMENT

In chapter 2, I made the case that we already have a large number of computer programs that are in some ways superhumanly intelligent, but they all have rather narrow capabilities. That chapter was presented in a mostly chronological fashion, as a history of superhuman AI. In this chapter, I revisit the historical development of AI but focus on the types of approaches used. I talk about how various approaches to developing artificial intelligence tried, and mostly failed, to achieve generality. The next two chapters will dig into two specific directions that are showing promise for achieving more general artificial intelligence right now.

Several of the superhuman AI systems discussed in chapter 2, including the Logic Theorist and Deep Blue, were products of what is generally called symbolic AI. In this approach, the knowledge is generally encoded as

propositional statements in a (somewhat) human-readable form.¹ In the Logic Theorist, the propositional knowledge included various axioms about numbers, as well as the shape of logical transformations. In Deep Blue the state of the board (e.g., “the white king is at A2”) is encoded propositionally, as well as the moves that are allowed (e.g., “pawns can capture diagonally forward”). A search algorithm is then used to try various recombinations of the existing statements, yielding new theorems in the Logic Theorist and plans in Deep Blue. This approach has many appealing features, such as that the knowledge base is human readable and human editable, we can guarantee that laws of logic are followed, and we can to some extent understand how our AI systems work. The symbolic approach to AI has produced plenty of research breakthroughs and useful systems. Services and apps you use every day, including search engines, map services, and databases, build on advances in symbolic AI.

The symbolic approach to AI is in some sense the original one, as it was espoused by most of the participants in the 1956 seminar that gave the field of artificial intelligence its name and articulated its first research vision.² Until around the 1990s or so, the symbolic approach dominated AI research. Its dominance was such that proponents of alternative approaches, who sought to infuse AI with new ideas and perspectives, came to refer to symbolic AI as “GOF AI,” or good old-fashioned AI.

Could we achieve AGI of some kind with purely symbolic AI? Alas, the evidence from multiple decades of trying is not encouraging. One problem is that encoding the knowledge the system needs in symbolic form is very hard. In the 1980s, many researchers attempted to create what was called “expert systems,” systems that would assist users with, for example, medical diagnosis or career advice.³ The idea was to interview human experts about how they did their jobs, and encode their knowledge and logic in the expert system. However, humans are not very good at articulating what they know or how they do their work in the crisp, unambiguous form that is required by symbolic systems. The problem got worse the more general the system was designed to be, something that does not bode well for creating general intelligence. A good illustration of this is the Cyc project, an ambitious project to encode all “commonsense knowledge” as propositional statements useful for symbolic AI.⁴ The project has been ongoing since 1984 and has had only limited practical results.

A related problem is how to interface with the “real” world beyond the symbolically encoded knowledge—in other words, how to give sensory inputs (vision, hearing, touch, etc.) to a system and let it affect the world by poking and tugging at it. The world’s first mobile robot, Shakey, was developed firmly within the symbolic paradigm.⁵ Shakey used a TV camera hooked up to a computer the size of a room (this was the early 1970s) to create an internal

representation of the world around it, and then planned its path and what actions to take using symbolic methods. The project was a huge success: not only did it furnish existence proof of a mobile robot, but several important algorithms were invented during the course of the project. However, Shakey was exceptionally brittle. It could only navigate and act within carefully lit and arranged rooms furnished with special-purpose blocks, and even then it could take minutes to decide on its next action. The sun shining through a window, or some dust floating past the camera, would upset all its plans. Unfortunately, this was a symptom of a larger malaise for symbolic AI: the world outside symbolically encoded knowledge turns out to be messy and complex, and not at all easy to encode in clean factual statements. When real progress was finally made in robotics in the 1990s, it occurred outside the symbolic paradigm by trying to mimic the functioning of insects.⁶

Researchers have made various attempts to philosophically understand the limitations that symbolic AI faces when it comes to expert systems or mobile robots. One such issue is the symbol-grounding problem: the symbols the computer system uses do not actually mean anything to it, because it cannot ground them in sensory experiences. To you, the word *fog* brings up memories of specific experiences and associations with moods and episodes; to a symbolic AI system, it is just a symbol that can be manipulated in specific ways.⁷

Starting in the 1980s, artificial intelligence research gradually focused more on machine learning. In machine learning, the AI system learns from data, or from its own behavior.⁸ While some early research into machine learning began in the 1950s, practical applications had been held back mainly by lack of data and computational power. Several algorithms developed in the 1980s, in combination with faster processors, made machine learning practical. In particular, as mentioned in chapter 2, effective methods for training neural networks were developed.⁹ But many other machine learning methods are not based on neural networks; in particular, methods based on decision trees and support vector machines continue to be important in real-world applications. Regardless of the algorithm used, you can roughly divide machine learning into three types: supervised learning, self-supervised learning, and reinforcement learning.

The most common form of machine learning is supervised learning. Here a neural network or another machine learning model is trained on sets of inputs and targets. The goal is to predict the target from the input. For example, the input can be a list of things a customer has bought, and the target is whether the customer is interested in a particular product category, such as male beauty products. By training on data from thousands of customers that have visited a store, the machine learning algorithm might learn to predict interests accurately so that the store

Starting in the 1980s, artificial intelligence research gradually focused more on machine learning. In machine learning, the AI system learns from data, or from its own behavior.

can target its advertising to particular customers. To take another example, the input can be an image of a face, and the tailoring can be to the age of the person in the image. In practice, supervised learning is used everywhere these days, from biometrics, where trained models predict if a fingerprint belongs to a given user, to health care, driving assistance systems in cars, credit decisions in banks, and grammar checking.

While early machine learning systems could only handle small or highly structured inputs, the deep learning revolution of the 2010s gave us supervised learning systems that operate directly on high-resolution images, as well as raw sound files and other high-volume unstructured data.¹⁰ AlexNet, discussed in chapter 2 as an example of a superhuman AI system, and the ImageNet data set it was trained on played a key part in this revolution. Interestingly, AlexNet itself was not very different from the neural networks that had been trained decades ago; it was just bigger and deeper (had more layers) and trained for longer on more data. The most important part of the deep learning revolution was arguably not new algorithms but more data and more computational power. The vastly larger amounts of data came via the internet, from millions of users and their digital cameras, and the increased computational capacities came from GPUs, graphics processors that were originally developed for video games but turned out to be really useful for training neural networks very quickly.

One way in which a deeper network can do things that a more shallow network cannot is by representation learning. Earlier machine learning systems, sometimes called “shallow” machine learning, needed their input data formatted into meaningful “features.” For example, say that you wanted to train a face recognition system that could identify individual people. In shallow machine learning, you would need to first use human-written software to identify a number of features of the face, such as distance between the eyes, mouth, and nose, the shape of the eyebrows, and so on.¹¹ With deep learning, you can do face recognition directly from the pixels of the image, without explicitly identifying the features. What makes this work is that the neural network learns its own feature extractors. The first few layers might identify lines and patterns of light and dark; in subsequent layers it might identify nose, eyes, ears, and so on. This process is called representation learning, meaning that the neural network learns the representations it needs itself.¹²

Supervised learning can also be used to train AI agents how to behave in games and simulations, or even real robots. Here the learning algorithm is used to imitate the actions of one or many humans. For example, let’s say you want to use supervised learning to teach a neural network to drive a car, based on data that were collected from human driving. The data would consist of lots of pairs of observations and actions. The input is an observation of the

state the car was in (camera data showing what is in front of the car, speedometer reading, perhaps other sensors such as lidar), and the target is the action the human took in that situation, such as turning the steering wheel left or braking. This works well.¹³ Perhaps surprisingly, the first demonstration of highway driving based only on supervised learning from human driving data occurred all the way back in the 1980s!¹⁴ The technology has kept getting better since then, and so has performance. Yet getting to driving that is consistently as good as a good human driver has proven elusive. Even the best of these trained neural nets sometimes make ridiculous mistakes.

Could we build AGI based purely on supervised learning? Almost certainly not. We don't have the training data. If we want to imitate our way to general intelligence, we need to have something to imitate first. And even if we could somehow log *everything* that a large number of humans did in an appropriate format for training, we would end up imitating all the mistakes, blunders, and bad decisions people make as well as the good ones.

Another form of machine learning is reinforcement learning. The roots of reinforcement learning go back almost as far as supervised learning: at least to 1952, when Arthur Samuels at IBM wrote a program that taught itself to play checkers by playing against itself.¹⁵ The core idea of reinforcement learning is that an agent learns to take appropriate action by observing the results of its own

actions.¹⁶ Every time the agent takes an action (such as making a move in checkers), it sees how the state of the world has changed (the new state of the board in checkers), and intermittently it also gets some kind of positive or negative reward. You may recognize this formalism from earlier in the book, as it is used for the definition of universal intelligence discussed in chapter 4. In reinforcement learning, the algorithm uses the feedback to update the neural network (here called the *policy*) so that it performs better. If the feedback was negative, it does less of what led to the negative reward, and if the feedback was positive, it reinforces whatever made it do what it did. At the beginning of reinforcement learning training, the policy takes random actions, but, as it learns from its mistakes, it gradually gets better.

While there were some early impressive applications of reinforcement learning, such as TD-Gammon, a famous backgammon-playing program from 1992, reinforcement learning saw few successes for decades.¹⁷ But then deep neural networks came along. In a breakthrough paper from 2015, a team of researchers at Google DeepMind showed how reinforcement learning could be used to train deep neural networks to play video games using only pixels from the screen as inputs.¹⁸ The algorithmic advancements were modest, but the researchers applied the algorithms to much larger neural networks than had been used previously, and trained them for much longer.

The result was neural networks that could surpass human performance on dozens of classical games from the Atari video game console. With newfound confidence, researchers started applying reinforcement learning to all kinds of things, from controlling simulated robots to designing houses and even the structure of other neural networks.¹⁹

So is reinforcement learning on its own a viable route to AGI? Unlikely. Applying reinforcement learning on its own requires that you can try out a task thousands or even millions of times; you won't find a good policy without failing again and again and again. This is fine for video games or for any task for which we have a fast and accurate simulation.²⁰ However, for most tasks in the physical world, we don't have that. You can't use reinforcement to learn how to drive a car in the real world, because you would likely have wrecked hundreds of cars and possibly killed people before the driving got any good. You can use reinforcement learning to drive in a video game like *Grand Theft Auto*, but despite what I thought before I got my driving license in Brooklyn at the age of forty-one, real driving is not like a video game. And building a sufficiently accurate simulation of driving, particularly city driving, is surprisingly hard.

Another big problem with reinforcement learning is overfitting, meaning that the policies tend to work only in situations very similar to those they were trained on.²¹ Say that you do manage to teach a neural network to drive. Chances are that the policy will only work in the particular

scenarios you trained it on: if you try to use this policy to drive on roads that are different from the ones you trained it on, it might do something entirely unexpected. It might also break down if you change the type of car or even just the weather. And don't even think of adding some unrecognized object (like a deer) that the neural network was not trained on.

For learning general intelligence in particular, we also have the issue of rewards. As you may remember from the discussion of universal intelligence, the assumption that everything that is worth doing can be formulated as a task that will give you a reward soonish is problematic. This is an issue in practice, not just in theory. You can define good reward schemes for many video games and most tasks in driving, as well as some robot tasks, such as trash picking. But, as discussed in chapter 4, many of the tasks we perform do not come with clear success criteria and therefore do not come with clear rewards. And while you could in principle define reward schemes for a large variety of tasks, including nice dance moves and going for a walk to clear your mind, doing this for everything we do in daily life would be a major undertaking, to say the least.

In biology, the only reward that counts in the end is evolutionary fitness: the propagation of your genes. And Darwinian evolution, essentially optimizing the “reward signal” of having lots of offspring, is what has produced the quite remarkable natural intelligence you are using

right now. Could we borrow this idea and create AGI evolutionarily? Like they do it on the Discovery Channel?

Evolutionary algorithms are algorithms that are inspired by natural evolution. They “evolve” solutions by making random changes to a population of solutions, discarding the worst ones and letting the best ones re-create, that is, making copies of them with small modifications. The solutions are evaluated by a “fitness function,” which is a numeric measure of the solution’s quality. By doing this repeatedly, the algorithm eventually finds good solutions. A solution here could be almost anything, depending on the task you want to solve. If we want to create intelligent agents, the solutions could be neural networks.

Could we create AGI through evolutionary computation? Intuitively, it would make sense that we could. We could create simulated environments where agents need to solve a large variety of tasks to survive, and the best agents (or rather the neural networks powering them) get to survive and have offspring. Eventually we should have highly and generally intelligent agents, if that is the winning strategy in that environment.

This sounds nice and easy, or at least I thought so back when I started my PhD. Yes, I was naive, befitting my age. Unsurprisingly (in hindsight), evolving general intelligence is easier said than done.

Evolutionary algorithms work well and are extremely robust when you define an appropriate fitness function.

Researchers have been able to evolve radar antennae, game levels, production schedules, and many other things, including neural networks that control mobile robots to successfully perform complicated tasks.²² However, in all these cases, the fitness functions are easy to define. General intelligence is, as we have seen, hard enough to define; writing a fitness function that somehow reflects it is harder still. (Yet another reason for wanting a good definition of general intelligence is that such a definition is the first step toward a fitness function.) In theory, the fitness function could be implicit, in that the agent performs better than other agents in a complex environment. In practice, such attempts have not so far led to the emergence of very intelligent agents, partly because simulated environments at the scale that seems to be necessary are very computationally expensive, and partly because it is hard to create environments with the right kind of challenges.

Now, you may wonder why we would need to use only one of these approaches to build AGI. Can't we combine them somehow? Get the best of each? And that's a good question. Throughout the history of AI there has been a tension between what have been called "neats" and "scruffies." Neats emphasize conceptually simple and theoretically well-founded solutions. An ideal AI system, to a neat, is based on a single principle, and it can be shown theoretically why it works. Scruffies, on the other hand, believe that anything goes, and argue that different approaches

to AI have different strengths, which can be usefully combined.²³ To a neat, an ideal AI system is a thing of beauty, preferably small and encapsulated in a few equations. To a scruffy, an ideal AI system is one that gets the job done, even if it's an unholy Frankensteinian combination of disparate parts and no one really understands how it works.

Some of the recently dominant work on deep learning is of the neat school, in that it is based on a single principle and algorithmically quite simple. For example, the AlexNet network that won the ImageNet competition is based on a simple principle (gradient descent) and only two different types of neural network layers. Using a modern deep learning framework, the code is only a few dozen lines. On the other hand, it is rather non-neat in that the sheer size of the trained network (sixty-two million parameters) makes it impossible to understand by simply looking at the network.

Other important recent work is decidedly more scruffy. To take another example of an extremely influential AI system, let's look at AlphaGo, DeepMind's program that in 2016 beat one of the world's best Go players and ushered in the era of superhuman play in this centuries-old board game.²⁴ AlphaGo is a symbolic AI system in the sense that the board state is represented symbolically, and a search algorithm that tests various moves sits at the heart of the whole system. However, it also relies heavily on supervised learning, which was used to train networks

Throughout the history of AI there has been a tension between what have been called “neats” and “scruffies.”

on thousands of games between master-level players. These networks help the search algorithm by evaluating both board states and possible moves. Finally, reinforcement learning plays a pivotal role in getting to that crucial human-beating performance. After the neural networks have gotten as good as they can by training on human data, they improve further by having the whole AlphaGo system play against itself, learning from its successes and failures.

In the next two chapters, we are going to look at two approaches to AI that build on and combine the various approaches we have seen here. The first is self-supervised learning, which is similar to supervised learning but predicts the data themselves instead of predicting targets. The other is open-ended learning, which combines reinforcement learning and evolution to invent new goals while training. These are the two approaches that I think have the best hopes for creating more general AI systems in the short or medium term. However, these approaches are very different, and so are the kind of AI systems they bring into the world.



SELF-SUPERVISED LEARNING OF FOUNDATION MODELS

Supervised learning, where we learn to predict targets or labels from inputs, is highly effective if you have enough data in the right format. For many tasks, you don't. For example, supervised learning works well for learning to identify human faces in videos, because it is easy to generate lots of training data by labeling existing videos with positions of faces. There are lots of tasks in the world for which such data sets exist, but unimaginably many more for which they don't. However, for many tasks, *unlabeled* data exist. For example, no labeled data set lets you use supervised learning to train a neural network to generate new, never-before-seen images. (By definition, the images cannot be labels if they don't exist yet.) However, the internet is full of images. Is there some way we can train a neural network on a huge pile of unlabeled images and make it somehow generalize, so that it can create new images?

Researchers have been working for decades on how to automatically create images, using a large variety of techniques. Some methods produced aesthetically interesting results but required a large amount of engineering efforts to get to a particular artistic style and were not flexible at all. Early attempts at training image-generators from data sets often used autoencoders, a type of neural network that can be thought of as hourglass shaped.¹ An autoencoder has a wide input layer, where the input can represent the whole image, and an equally wide output layer, with smaller layers in between. When training the autoencoder, you simply ask it to predict its own input. In other words, you input images and reward it for producing exactly the same images as outputs. In supervised learning terms, you use the input as a label. (This is why it's called *self-supervised*.) Now, this would be trivial if all the network layers were the same size. But because of the hourglass shape of the autoencoder, where information must flow through a bottleneck middle layer, the network must learn to compress the input by focusing on the most important information. You can think of it as pushing the image through a pipe that is too small for it, so that only the most important parts of the image squeeze through the bottleneck. If you train an autoencoder on a few tens of thousands of images, it would learn features that are general to all those images. You could then input almost anything into the bottleneck layer of the autoencoder and

get a nice image as output. At least in theory. In practice, the results of these early autoencoder experiments were not so impressive.

In 2014, a new self-supervised method entered the stage and revolutionized image generation: the generative adversarial network, or GAN.² A GAN consists of two separate neural networks, the *generator* and the *discriminator*. The generator and the discriminator are engaged in a competition of sorts: the discriminator is trying to learn to tell real images apart from the fake images made by the generator, and the generator is trying to learn to fool the discriminator. To train a GAN, you need a large number of images (usually at least a few thousand, but often a few million are used). Initially the generator generates random images, and the discriminator makes random judgments, but over the course of generating and judging lots of images lots of times, they gradually get better. In the first cycles of training, the generator learns to produce noise that is just sufficiently image-like to fool the discriminator, after the discriminator learns to spot the generator's feeble attempts at deception; then the generator steps up its attempts, but so does the discriminator. It is instructive to see the evolution of image quality as a GAN is being trained, going from noise to weird blobs to fuzzy images to, eventually, if all goes well, high-quality images.

Within a few years after their invention, GANs were capable of generating photorealistic images at high resolu-

tion in a large variety of styles.³ People were beginning to take note, and using GAN-generated images in place of photos of certain things, notably human faces, became feasible. However, the method is still somewhat unreliable. In particular, it's hard to tell a GAN what to create; or rather, it doesn't listen very well.

The next big advance in self-supervised learning for images came along soon afterward. Diffusion models became famous in 2021 with OpenAI's Dall-E system (although they were originally invented in 2015 and had been developed by various researchers since).⁴ A diffusion model is, like an autoencoder or a GAN, trained on a very large number of images (the more, the better). But the training works differently. In a diffusion model, noise is added to each image step by step until the image consists only of noise. Then the network is trained step by step to remove the noise. So it is trained to reconstruct an image that is 49 percent noise from one that is 50 percent noise, and then one that is 48 percent noise from one that is 49 percent noise, and so on until zero. After training on reconstructing lots of images from partial noise, the network is able to construct completely new images from complete noise. It has learned the general pattern of the images it has been shown, and can use this knowledge to create completely new images in the style of the images it has been shown. Compared to GANs, the process is slower but more stable.

All of these methods can also be taught to generate images in particular styles, or with particular motives. This is done by associating images with their captions during training, so that the generation of a new image can be guided by writing a new caption.

As I'm writing this, image generators trained with self-supervised learning have become so capable that they can be used for a number of practical purposes. I myself use them to illustrate slides in my presentations, and I know people who use them for generic illustrations for marketing, educational materials, and other purposes. There have also been high-profile cases of believable fake images of recognizable people, such as a very popular one of Pope Francis wearing a stylish coat from the fashion brand Balenciaga. The quality of the generated images is constantly improving, and we have recently seen the first few decent-quality generated videos. There are many good image-generating networks, such as Midjourney, developed by a small company of a dozen or so employees, and Stable Diffusion, which is open-source.⁵ New versions of these models are released at a steady pace.

Although images are important media, and the consequences of high-quality automated image generation can be huge, it is in the text domain that generative AI based on self-supervised learning has led to particularly impressive advances—so advanced that some have started talking about self-supervised learning on internet text as a path to AGI.⁶

Self-supervised learning on text has a long history. Claude Shannon, inventor of information theory and a crucial figure in the development of early computers, was probably the first to propose a form of self-supervised learning.⁷ To understand what he did, first consider what would happen if you tried to “write” English text by randomly selecting letters between A and Z, where each letter would have the same probability:

XFOML RXKHRJFFJUJ ZLPWCFWKCYJ
FFJEYVKCQSGHYD QPAAMKBZAACIBZL HJQD

In 1948, Shannon measured the frequency of letters in the English language (well, actually in a few books) and produced a table of how common the various letters are. The most frequent letter is *E* and the least frequent is *Z*. Once you have such a table, you can “sample” from it: you can select letters randomly, proportional to how common the letter is. Doing this, you can get sequences of letters that are somewhat English-like, even though they are not quite words. At least, the resulting text no longer looks like Klingon:

OCRO HLI RGWR NMIELWIS EU LL NBNESEBYA
TH EEI ALHENHTTPA OOBTTVA NAH BRL

But we can do better than this. In particular, we can make a table of not only how frequent letters are on their

own, but how frequent they are depending on which letter comes just before. For example, *k* is unlikely to follow *t* but quite commonly follows *c*. This is called *conditional* probability, where the previous letter is the condition. Sampling from a table of conditional probabilities based on the English language, where every letter's likelihood depends on the previous letter, we get strings like the following:

ON IE ANTSOUTINYS ARE T INCTORE ST BE
S DEAMY ACHIN D ILONASIVE TUCOOWE AT
TEASONARE FUSO TIZIN ANDY TOBE SEACE
CTISBE

By now it's probably clear to you how we can do even better: by making every letter's probability conditional not just on the previous letter but on the two previous letters. Or why not the three previous letters? Shannon called this kind of structure an *n-gram*, where a 0-gram is a table of unconditional letter probabilities, a 1-gram is a table of probabilities conditional on the most recent letter, a 2-gram means the probability depends on the two most recent letters, and so on. Here's some text produced by a 3-gram trained on English text:

IN NO IST LAT WHEY CRATICT FROURE BIRS
GROCID PONDENOME OF DEMONSTURES OF
THE REPTAGIN IS REGOACTIONA OF CRE

As you can see, the output gets better and better. The 3-gram includes some actual English words, and some sequences of letters that would not surprise anyone if they turned out to be used as words in some obscure context. If I went to a doctor, and the doctor told me I had a grocid pondenome, I would certainly be worried.

You can also use n-grams on the word level. Here are words sampled according to their probabilities in the English language:

REPRESENTING AND SPEEDILY IS AN GOOD APT
OR COME CAN DIFFERENT NATURAL HERE HE
THE A IN CAME THE TO OF TO EXPERT GRAY
COME TO FURNISHES THE LINE MESSAGE HAD
BE THESE

And here each word is sampled based on its probability conditional on the previous word:

THE HEAD AND IN FRONTAL ATTACK ON AN
ENGLISH WRITER THAT THE CHARACTER OF
THIS POINT IS THEREFORE ANOTHER METHOD
FOR THE LETTERS THAT THE TIME OF WHO
EVER TOLD THE PROBLEM FOR AN UNEXPECTED

Despite their simplicity, n-grams can be surprisingly powerful. Unfortunately, they are limited by the amount

of data they need. More text is needed to calculate the conditional probability tables as n gets larger. The tables themselves also get very large. This is because the more letters you depend on, the more rare that combination of letters is. In fact, the amount of text you need to calculate n -grams grows exponentially with n , meaning that this method only works for small n . Still, small n -grams can be very useful. If you had a mobile phone in the 1990s or early 2000s and remember the system that helped you write text messages using only your numerical keypad, that was an n -gram.

Could we not just use a neural network instead, give it the n previous letters, and train it to predict the next letter? In principle, this should allow us to learn more general patterns. Neural networks are generally very good at interpolating between points in the training data. In this context, you could hope the network can make a reasonable guess at which letter comes next even if it has never encountered the exact combination of preceding letters.

As neural networks became popular again in the 1980s, various neural network approaches to text generation were explored.⁸ The first attempts involved simply feeding the n preceding letters (or words) into the network and predicting the next letter (or word), just as you would do with an n -gram, but in this case you train a network instead of updating cells in a table. However, this approach has the limitation that the number of previous letters you

depend on is fixed and can't be very long or else you would need to train a giant network. In effect, these simple networks have extremely limited memory, and by the end of a sentence, they could easily forget what the beginning of the sentence was about.

To overcome this limitation, recurrent neural networks were invented.⁹ These networks are built and trained to be run not once but many times in sequence. Each time a recurrent neural network is run, it uses its own internal state from the previous time step as an input. This way, it can retain memory of what it did in the past. Using a recurrent neural network to produce text typically proceeds by first using a short snippet of text as input (the “prompt”), and then using the network to predict the next letter (or word) and adding it to the prompt. Then you use it to produce the next letter or word, based on the current state of the network (which is based on all the text it has seen so far) *and* the newly outputted letter or word. In other words, the network is using the text it outputs as parts of its input. These networks are called recurrent because the output *recurs*.

Although the basics of recurrent networks were established in the early 1990s, it took the advent of the World Wide Web, which made unprecedented amounts of text freely available, to really make them work. It also took some technical innovations; in particular, the long short-term memory (LSTM) network, which was introduced in 1999, uses a number of clever ideas to make it

easier for these networks to form long-term memories.¹⁰ And, of course, we also needed better hardware to train these complicated neural networks on large amounts of data without having to wait years for the training runs to finish. In particular, working out how to use GPUs (graphics cards) for training was key to the success of recurrent neural networks.

How well do recurrent neural networks work? Very well, as long as the text is not too long. LSTM networks in particular are used in tons of places for generating, completing, spell- and grammar-checking, and classifying text. Popular email services, including ones you have almost certainly used yourself, rely on LSTM networks to help you compose your messages. An LSTM network trained on a large amount of text is typically capable of generating a paragraph or so of coherent text in response to a prompt. For example, I just fed the following prompt to a large pretrained LSTM network: “The book Artificial General Intelligence, written by Julian Togelius”. Here is the full paragraph the network gave in return:

The book Artificial General Intelligence, written by Julian Togelius, was published in 1985 and was published by Science, the United States, and at the Harvard Business School. The book is full of illustrations and insights to the topic of artificial general intelligence.

Factual errors aside—I did not publish any books when I was six years old; in fact I could barely read books in Swedish, and certainly not in a foreign language such as English—the text is more or less coherent and plausible. We find at least one grammatical oddity (you don’t publish books “at” a business school) but you would probably just think it was a typo if a human had written the text. One also does not typically refer to a publisher (“Science”) by country (“the United States”). Note how the final words of the paragraph (“artificial general intelligence”) recur from the prompt, showing that the system has some memory. But where do the wrong yet oddly specific references to 1985 and Harvard Business School come from? The answer is that these are just likely continuations of the text the network has seen so far. The network is not consulting a database of true facts about the world; it is just doing the same thing as Shannon’s n-grams: outputting the most likely next letter. It’s just doing it via a much larger and more complex structure; in this case, the trained LSTM network has many millions of parameters. And when you think about it, 1985 is not an unlikely year for a book about AI to be published (much more so than any year before 1960), and Harvard Business School is the kind of place where people talk about AI.

The next advancement in self-supervised learning for text can be traced to a paper from 2017 with the curious title “Attention Is All You Need.”¹¹ That paper introduces

the transformer, a type of neural network based on self-attention. The “attention” here refers to a neural network layer that uses its input to select which of the outputs of the previous layers to focus on. What this means in practice is that the transformer can more easily learn to consider specific parts of longer texts when generating its output. It is also easier to scale transformers to very large network sizes. While the transformer architecture is novel, it is still a neural network, and it is still just trying to output the most probable next token.

Since the transformer architecture was invented, multiple teams around the world have competed in designing and training ever larger networks on ever more data scraped from the web. Because they are so large, text predictors based on transformers are generally called large language models, or LLMs for short. For efficiency, LLMs usually predict “tokens,” where a token is either a character, a short word, or part of a long word.

In 2019, the AI research company OpenAI unveiled Generative Pre-trained Transformer 2, or GPT-2 for short, a network with 1.5 billion parameters trained on 40 gigabytes of text.¹² GPT-2 caused a stir in AI circles by being able to generate high-quality text in a variety of styles, including newspaper articles, scientific reports, and nursery rhymes. The output of GPT-2 is generally coherent up to about a page or so of text, reflecting its context length of 1,024 tokens. In a move that was extremely beneficial to

AI research, OpenAI released the trained GPT-2 freely for anyone to use and modify, leading to an explosion of activity in modifying and fine-tuning the model.

Here it makes sense to explain the term *foundation models* used in the chapter title. Although the term is not loved by all, a foundation model is a large machine learning model trained for a long time on a lot of data, which can be used for multiple tasks. Often this includes tasks the model was not originally trained for. For example, the Stable Diffusion model was trained to generate images based on text embeddings, but the same trained network can be used to “inpaint” missing parts of existing images, to restore images, to upscale or downscale images, and even to translate between visual styles. Similarly, large language models such as GPT-2 might only be trained to generate text as continuations of prompts, but they can be used for a whole range of tasks, including compressing text, translating between languages, and judging the similarity of text snippets well. (More on the surprisingly varied capabilities of LLMs soon.)

Another thing that you can do with foundation models is to fine-tune them. A trained model is also referred to as *pretrained*, meaning that it has been trained for a number of things, but not the specific data set you may want to train on. But the pretraining has enabled the neural network to learn a number of useful representations, so that it can better use the new data. In a recent project of ours, we

A foundation model is a large machine learning model trained for a long time on a lot of data, which can be used for multiple tasks.

took an LLM that had been pretrained on large amounts of text from the internet but not, presumably, game levels. We then fine-tuned it on a set of levels for the puzzle video game *Sokoban*, converted to text form. The fine-tuned network could then generate new *Sokoban* levels, something it could not do before fine-tuning. Networks that were pretrained learned faster and better than networks that were untrained. Apparently, training on text somehow taught the network useful representations that allowed it to learn to generate game levels, even though game levels do not seem similar to blog posts, newspaper articles, books, and whatever else the LLM was originally trained on. This phenomenon, where pretraining helps a neural network learn when fine-tuned on other data, is very general and part of the reason foundation models are so useful.

Although GPT-2 made a big splash among AI researchers, especially those interested in natural language understanding, it took a few more developments for people outside the profession to notice that something special was afoot. In 2020, OpenAI released GPT-3, a giant neural network with the same structure as GPT-2, but about a hundred times larger.¹³ GPT-3 is capable of generating better text that stays coherent for longer and also seems to know much more about the world—not unsurprisingly, as GPT-3 is much larger and trained on much more data. With clever prompting, you can make GPT-3 write multipage essays that are quite good on a large variety of

topics, poems of varying quality, and also dictionary-style descriptions that are often mostly correct. However, it is all in the prompting. Prompted with “The 1998 invasion of Sweden,” GPT-3 gives us the following:

The 1998 invasion of Sweden by NATO forces was a military operation that took place on October 17, 1998. The invasion was a response to a Swedish request for assistance after a terrorist attack on the Swedish parliament. NATO forces, led by the United States, invaded Sweden and quickly secured the country. The invasion was largely successful, and Sweden was able to quickly restore order. However, the invasion was met with criticism by some Swedes, who saw it as a violation of their country’s sovereignty.

This is very believable in every way—except that it is not true. It is merely likely, meaning that it is a likely continuation of the prompt. (Although if it did happen, I might have missed it; it was a Saturday, and I was nineteen years old, so I probably had a hangover.)

During the next two years, GPT-3 continued to improve in various ways as it was trained on more data; in particular, it became better at programming. Prompts like “this is a Java method that takes a string as input and returns the number of consonants in it” are likely to result

in useful code. But it was when ChatGPT, a descendant of GPT-3, was unveiled in December 2022 that the world really sat up and took notice. ChatGPT is not much more advanced in terms of what text it can generate, but it is much easier to prompt, as you can simply tell it what to do. The prompt “write me a limerick on the theme of sentient rhubarbs” results in exactly that. It even inserts a not-terrible pun:

There once was a rhubarb so bright // It gained
consciousness overnight // It could talk and it walked
// And it loved to be stalked // This sentient veggie
was quite a sight!

One of the main innovations incorporated into ChatGPT is instruction tuning. This technique, invented by a team at Google, is quite simple.¹⁴ You instruction-tune an LLM by additionally training it on many thousands of texts that take the form of instructions and responses. The instructions come in many different forms; some are questions, and many are imperatives, such as

Please answer the following question: What is the melting point of steel?

Tungsten melts at 3422 C. This is one of the highest melting points of any metal.

By fine-tuning on large and diverse sets of instructions and their answers, an LLM learns the general pattern of instruction following. While it would seem that some extra fine-tuning would add little to the capabilities of an already trained LLM, in practice it turns out that instruction tuning makes LLMs much easier to interact with. It unlocks a more natural, conversation-like way of controlling these capable but mysterious beasts.

The other main innovation (again invented elsewhere) incorporated into ChatGPT is reinforcement learning from human feedback (RLHF).¹⁵ The human feedback in RLHF consists in humans selecting the best of multiple completions of the same prompt. A separate model is trained to replicate the humans' preferences, and that model is then used to iteratively adjust the LLM using reinforcement learning. It turns out that RLHF is very useful for "taming" a language model. ChatGPT is quite good at not explaining how to mix explosives, bully people, or commit suicide, largely because of the RLHF procedure, where nicer and more rule-abiding responses have been reinforced.

The release of ChatGPT in December 2022 triggered a huge public interest in LLM technology, as it was so easy for anyone with an internet connection to use the model, and so easy to make it do interesting things through the chat interface. Students all over the world experimented with letting ChatGPT do their homework, software developers

experimented with letting it write code for them, professors asked it to create exam questions, and people asked it to create shopping lists, song lyrics, and recipes.

Almost everyone who has used the system has also tried asking it straightforward factual questions, often with mixed results. ChatGPT is still prone to hallucination, like other LLMs, and will happily make up facts. While it rarely gets obvious facts wrong (the capital of France is still Paris), it just today told me that two esteemed AI researchers I know personally are married to each other. They are not.

In March 2023, OpenAI released GPT-4, a larger and more capable model along the same lines as the previous GPT iterations incorporating the advancements of ChatGPT.¹⁶ While OpenAI will not tell us how large exactly GPT-4 is, nor what it was trained on, it seems clear that it has been trained on a larger data set. Generally, it can be seen as a more capable version of ChatGPT. A large advancement is that it has a longer context length, so that it can understand and generate longer texts.

LLMs as such are not agents. They don't do anything of their own volition. They just sit there, waiting for your prompt and producing some text in response. However, it is possible to build agents out of LLMs. In a computer, everything is fundamentally a string of characters, which is exactly what an LLM takes as input and produces an output. Communication between different programs

happens by sending text back and forth. You can easily build a program that feeds the output of an LLM into another program and then feeds the output back into the LLM, together with some prompt. You could prompt the LLM to write code for something, run the code in a program interpreter, and return the results to the LLM. You can also simply ask an LLM to reason about something, and feed the results of that reasoning back into the LLM. A number of experiments in this vein are ongoing as I write this. A team at Google has developed SayCan, a robot control system that feeds the LLM descriptions of the world around it, asks it to reason about them, and parses its output to control the robot.¹⁷ An open-source project called AutoGPT asks you to give it a high-level goal, and then asks GPT-4 to break this goal down into subgoals.¹⁸ AutoGPT keeps a list of the subgoals, gives them one by one to GPT-4, and asks it to carry them out using web searches, code execution, and other tools.

Another thing that LLMs can easily be combined with is image generators, the other big class of foundation models we talked about at the beginning of the chapter. GPT-4, and certain other LLMs, can both analyze and create images. You can feed an image into the LLM and ask it to analyze it, then ask it to generate new images in the same style or make edits to the original image.

While OpenAI has designed and trained the most famous LLMs, other models of similar capability have been

trained by companies such as Google, Mistral, and Meta and are freely available. This has led to a boom in open-source activity where companies, academic researchers, and even individual hobbyists fine-tune various models, build applications on top of them, and make them available on the web for everyone to use. Keeping up with the breakneck pace of developments in this space is difficult, but we can safely assume that whatever the frontier capabilities of LLMs are right now, they will soon be available in an open-source model.

The enormous versatility of modern LLMs is impressive, and they will very likely be ubiquitous in the future. The ability to combine LLMs with other software allows them to do things that they couldn't do otherwise, such as web searching, arbitrary code execution, or even just arithmetic. Still, LLMs have some glaring shortcomings. Perhaps the biggest one is the apparent inability to consistently stick to the truth. In this context, it is worth remembering that LLMs were never trained to be factual; they were trained only to predict the next token, just like an n-gram. We do not currently know whether we can ever make LLMs consistently truthful. Other issues include the tendency of LLMs to replicate the biases inherent in the texts they were trained on, including gender and racial biases (more on this later). LLMs are also very bad at some simple computational tasks that can easily be done by other software, such as planning, pathfinding, and

even standard multiplication. All of this is compounded by our lack of real understanding about how LLMs and other foundation models work. We understand the neural network architectures because we designed them, and we understand the training procedure, but we don't understand the internal representations that have been learned.

One way to illustrate the uneven capabilities of modern LLMs is to ask them to play chess. It turns out that GPT-4 knows a whole lot about chess and will often suggest very clever moves worthy of master players. But it can also frequently suggest illegal moves, as if it doesn't even understand the basic rules of the game. To complicate matters further, GPT-4's skill at chess varies greatly depending on how, exactly, you prompt it. It seems possible to get GPT-4 playing chess at a reasonably high level (though not as high as Deep Blue, not to mention modern chess engines), but that requires a skilled human operator prompting it in the right way. This raises the issue of determining how much of the skill is due to GPT-4 and how much is due to prompting it well; prompting LLMs well is a deep and arcane art. Fascinatingly, mentioning a famous chess player, such as Magnus Carlsen, seems to produce better results. If you want a bot to play a good game of chess, you're much better off with a special-purpose chess engine than with an LLM.

When you think about it, it is amazing how many things LLMs can do without being explicitly trained to

The ability to combine LLMs with other software allows them to do things that they couldn't do otherwise, such as web searching, arbitrary code execution, or even just arithmetic.

do them. Remember, they are just trained to predict the next token. One way of thinking about what LLMs do is that they role-play. LLMs try to complete the text you have written in the style of the text you just wrote. Imagine that you were doing the same thing. For example, someone gives you the first two sentences of an imaginary speech by Winston Churchill and tells you to finish it the best you can. You would probably use real names and events where you could, and invent plausible-sounding ones otherwise. Applications built on top of LLMs (such as the chat mode of Microsoft's Bing search engine) generally feature complex prompts that instruct the LLM as to what role it should play.

Another way of thinking about what LLMs do is that they learn algorithms for predicting the world. We have long known that a recurrent neural network could in principle learn any algorithm. Until recently, this was mostly thought of as a theoretical possibility, as learning nontrivial algorithms with neural networks seemed far too computationally complex to be practical. However, for predicting a vast amount of diverse text, maybe the easiest thing for the network to do is to learn general algorithms.

Will LLMs and other foundation models eventually lead to AGI? Who knows. Some researchers argue that GPT-4 already shows "sparks of AGI" given the wide range of things it can do.¹⁹ Others would say that the problems with truthfulness will not be overcome, and it will not

be possible to build LLM-based agents that are reliable enough to be widely used.

There is also the lurking suspicion that LLMs, which are after all only trained to reproduce text that is already written, are limited in how much new knowledge they can produce. After interacting for a while with LLMs, you start to recognize the genericness of their responses. Even when an LLM produces what might count as new knowledge, for example, by applying an existing methodology to a new data set, you might argue that this is merely pattern matching. Researchers have learned that LLMs are much better at solving versions of problems that are already found on the internet; for example, an LLM performs much better at solving common ciphers than novel, similarly complex ciphers.²⁰ So while chatting to an advanced LLM such as GPT-4 sometimes feels like having an intelligent conversation, at other times it seems like you are engaging with a clever interface to a low-quality, approximate database drawn from random corners of the internet. Could we build AGI some other way, without focusing on replicating what humans have already done? Perhaps. In the next chapter, we will explore attempts to foster open-ended learning in artificial agents.

OPEN-ENDED LEARNING IN VIRTUAL WORLDS

We humans are probably the most general intelligences on this planet (though, of course, this depends on the exact definitions of *general* and *intelligence*). It stands to reason that we should be able to create other general intelligences by using the same algorithm that created us. That algorithm is almost certainly Darwinian evolution, first on a purely genetic level and later on also on a cultural or memetic level.¹

As described earlier, evolutionary computation is an attempt to use Darwinian natural selection as a method for finding solutions or generating programs.² And evolutionary algorithms really do work well for many well-defined problems. For less well-specified problems, such as creating artificial general intelligence, it is difficult to come up with a good fitness function. A closely related problem besets reinforcement learning, which is the other

approach to learning from taking actions in the world: it is very hard to design a good reward function, beyond learning to perform individual tasks. How do you reward general intelligence if you can't even define it?

In the case of natural evolution, it took billions of years to reach *Homo sapiens sapiens* (you and me) through the slow process of letting whatever genes that could create organisms that survived and reproduced prevail. The environment we evolved in was hugely complex and constantly changing, and evolution took anything but a straight path. Our anatomy is full of evidence of how organs and body parts that evolved for one thing were co-opted into another thing. Did you know that feathers probably originally evolved for insulation rather than flying? Once some reptilian species were covered in feathers, they discovered that their plumage allowed them to glide. Those reptilians that could glide better by flapping their forearms had an evolutionary advantage, and soon (well, after a few million years) their offspring were birds.

Open-ended learning is an umbrella term for various approaches to cocreating agents and environments where agents can learn or evolve over a long time, gradually developing richer and more complex behavior repertoires. Generally, this means that we are focusing not only on how a single agent learns but on how multiple agents learn together, or on the generation of problems or environments that they can learn from. This is a radically

different approach to AGI than learning from words and images created by humans.

As with most things in AI, the idea of open-ended learning has been around for a long time, waiting for the computing power and software infrastructure that would allow it to be fully realized. We are still waiting. However, more and more developments in the space of open-ended learning suggest that we might soon see interesting capabilities emerge. Most of the early attempts at achieving open-ended learning took place within the loosely structured research field of artificial life, which aims to study "life as it could be" as opposed to life as it is (the domain of biology).³ Artificial life generally takes a broader view than artificial intelligence, reasoning that intelligence is just one facet of biological life and may not be separable from the organism that exhibits it. Much work in artificial life therefore tries to create simulated beings that exhibit some aspects of life, which may include intelligence.⁴

Early progenitors of open-ended evolution in an artificial life context include Nils Aall Barricelli's experiments in simulating evolution on an experimental computer in Princeton in 1953.⁵ One of the more sophisticated early attempts is the Tierra system, a simulation first run in 1990, where various programs inhabit a shared memory space and compete for processor time and space.⁶ They are able to use their limited time to make copies of themselves and change or overwrite other programs. The

Artificial life generally takes a broader view than artificial intelligence, reasoning that intelligence is just one facet of biological life and may not be separable from the organism.

competition gives rise to an evolutionary process where programs evolve to dominate the space, often developing dirty tricks along the way. For example, some programs evolved to become parasites on other programs, mimicking parasite-host relationships in nature.

Many attempts to produce open-ended evolution rely on algorithms trying to mimic cooperative or competitive coevolution. Coevolution is a concept from biology signifying that two species evolve dependently on each other, because the organisms of the two species interact in some way.⁷ Many forms of coevolution exist. The coevolution of various species of insects and flowering plants provides good examples of cooperative coevolution, where the fitnesses of members of different species depend *positively* on each other. In other words, they have a win-win arrangement. A common pattern is that flowering plants produce nectar, which insects eat for nourishment; at the same time, the flowers' pollen sticks to the insects and is transported by them to other plants of the same species. Thus the insects and the plants form a symbiotic relationship. The insects get food, and the flowers get to spread their pollen. The coevolutionary dynamics between specific pairs of species have led to the evolution of complex functionalities that would likely not have evolved otherwise. For example, some plants have evolved pigments for particular colors, and their pollinator insects have evolved eyes capable of identifying these colors. Other plants

attract their insects with scents, or even flowers that look like the insects in question. Much of the beauty of an early summer day in the garden results directly from cooperative coevolution.

Whereas in cooperative coevolution the individuals of both species gain in fitness, competitive coevolution describes coevolutionary scenarios where one individual's gain is another individual's loss. Predator-prey dynamics are typical examples of competitive coevolution. Foxes and rabbits have an obvious such relationship, the latter being prospective food for the former. Gradually, rabbits have evolved oversized legs, permitting a quick getaway; eyes placed to cover an almost 360-degree field of view; and strategies of suddenly and unpredictably changing direction while hopping. Foxes have evolved fast running gaits, accurate eyes offering stereo vision with depth estimation at the cost of a narrow field of view, and strategies that involve ambushing their prey. Higher fitness for a particular fox will likely come at the expense of lower fitness for a particular rabbit, and vice versa.

Much of the astounding variety and inventiveness of the natural world comes down partly to coevolution. Could we replicate such dynamics in computer code? It would seem to be a major advancement for evolutionary search to not have to manually specify a fitness function.

In an influential experiment from 1990, Daniel Hillis showed that coevolution can be practically useful.⁸ Hillis

used artificial evolution to evolve sorting networks, simple circuits that can sort numbers. To assign a fitness to each network, he first tested them on several sorting problems. However, this gave rise to networks that were specialized to the particular problems on which they were tested. He then started evolving the sorting problems as well as the networks. The networks were, as before, rewarded for solving as many problems as possible. The sorting problems were rewarded for fooling (being unsolvable by) as many networks as possible. In the way Hillis formulated it, the sorting problems were “parasites” on the sorting networks—parasites and hosts being another common type of competitive coevolution in nature. This turned out to work very well, and the coevolutionary process resulted in better sorting networks than standard, fixed-objective evolutionary algorithms could find.

In the 1990s and early 2000s, researchers tried getting competitive coevolution to work in robotics, either using miniature real robots or simulated robots. In particular, many researchers tried to replicate predator-prey dynamics, with one robot being the predator and the other serving as prey.⁹ The robots would keep chasing and evading each other again and again, getting better at outsmarting each other as their neural networks evolved. Some of these experiments showed interesting parallels to real-world evolution. For example, in simulated environments where the positions of the robots’ “eyes” (vision sensors) were

also evolved, the predator robot would evolve forward-looking eyes with capability for depth vision (like a fox), and the prey robot would evolve side-looking eyes so that it could not be ambushed (like a rabbit).¹⁰

Still, these experiments in coevolution have by and large failed to evolve complex strategies. One common phenomenon is "cycling," where the same strategies evolve over and over. This works a little bit like the game rock-paper-scissors, where paper beats rock, rock beats scissors, and scissors beats paper. In the same way, the predator could evolve a strategy that is beaten by a new strategy from the prey, which in turn is beaten by the previous strategy from the predator.¹¹

Why did these various experiments in artificial coevolution fail to create complex, impressive intelligence? Well, coevolution in nature happens in a complex world, rich in sensory impressions, with near-endless possibilities for new strategies. It is more than likely that these experiments were just too simple. The robots (simulated or real) were too simple: too simple bodies, too simple actuators, too simple sensors. Above all, the environments were just too simple. How could you possibly learn general intelligence in a (simulated) barren room with nothing to do except chase another cylindrical robot that always moves at the same speed?

So we would need more complex environments. Where could we get them? Video games, that's where.¹² Modern

video games have huge worlds full of interesting things to learn and do. They are also designed so that humans can learn to play them.¹³ In fact, a leading theory of game design says that a main reason why we find games fun to play is that we get better at them, and successful games are at their core well-designed pedagogical experiences.¹⁴ The wide variety of video game designs is a testament to game designers exploring new and interesting challenges they can give to the human brain. Some of the most popular video games have a rather extreme skill depth; people can play *StarCraft* or *Street Fighter* for decades and still keep getting better at them.¹⁵ Could we simply let loose our learning algorithms on video games and let them learn ever more complex skills and ever more general intelligence?

AI in video games is what I do all day, every day, so I will have to stop myself from devoting the rest of the book to this topic.

I discussed video games in chapter 6, in the context of deep reinforcement learning—in particular, the 2015 paper that ignited massive interest in combining reinforcement learning with deep neural networks. That paper described training networks on a few dozen games from the classic Atari 2600 game console from 1976. Yes, that's the old box with a wooden front that connected to your (color or black-and-white) TV. It's an extremely limited device by any standard—your modern cell phone has millions of times as much memory—but it can run a bunch

of decent games and was a big success back in its day. You might have played versions of some of those games, such as *Pac-Man*, *Space Invaders*, *Donkey Kong*, or *Q*bert*. These days, the Atari games can be emulated very quickly on modern computers, which is why they are so useful for reinforcement learning, which generally requires tens of thousands of playthroughs to learn to play a game. In 2015, it was a big sensation that neural networks could be trained to play these games from visual inputs only.¹⁶

If you recall the discussion in chapter 6, the problem with training deep neural networks to play games is that they tend to “overfit,” or specialize too much. A neural network trained on a single Atari game cannot play another game. In fact, it can usually not even play levels *of the same game* it has not been trained on already. Additionally, adding some noise to the visual input or rescaling the resolution would also completely derail the trained network.¹⁷ The “intelligence” displayed by these reinforcement-trained networks seems very brittle, the opposite of general.

In 2018 and 2019, several different research groups tried to analyze this phenomenon further and develop ways to train deep networks to have more robust game-playing capabilities.¹⁸ One of these groups was my lab at NYU. We developed versions of a few simple 2D video games (including *Frogger* and *Boulder Dash*, in case these names ring any bells for readers with a taste for 1980s

games) where we could easily change the levels and various other parameters.¹⁹ We confirmed that by training on individual levels of these games, standard deep reinforcement learning algorithms indeed learned to play only the particular levels on which they were trained; give them another level of the same game, and the trained networks would just perform meaningless actions, like trying to walk into a wall. If you trained on multiple levels, performance got a little better. But not much better.

We then started generating levels. We created a system that constantly provided new levels for training the agent.²⁰ In every episode (attempt at playing the game), the deep learning agent played a completely new level: walls, enemies, keys, and doors appeared in new places every time. The generator also had a “knob” for difficulty: by specifying a higher difficulty level, the generator would create more complicated mazes and more enemies. When the agent had learned to finish most levels of a particular difficulty, the system would increase the difficulty, similar to how a teacher gives you more advanced tasks as you gain mastery of a skill. This worked—at least in the sense that the agents trained on an endless supply of machine-generated levels were much better at playing levels they had never seen before than were the agents trained on just a few human-made levels. Diversity breeds robustness.

But there were limits to what these agents could do. They never learned to finish the most difficult levels, and

of course they never learned to play any games other than the one they had trained on. Could we take this idea further? Could we generate a wider range of levels and game variations, perhaps completely new games?

Some video games have incorporated level generation for decades.²¹ One such classic game is *Rogue*, a complex dungeon-crawling game where players face a completely new dungeon every time they play. *Rogue* was so influential that it gave a name to a whole genre of games, *rogue-likes*, which all prominently feature the generation of environments or other aspects of play. *Spelunky*, *Hades*, and the *Diablo* series are prominent examples of more modern roguelikes. More recently, we have seen games like *No Man's Sky*, a gigantic space exploration game that features quintillions (!) of planets, all with unique geology, flora, and fauna.

Researchers have also been exploring how to generate completely new games, although this is much harder than generating levels, quests, or items for existing games. A few successful board games have been AI generated, including the game *Yavalath*. *Yavalath* was not designed by the noted game designer and researcher Cameron Browne; instead it was designed by Ludi, an AI system that Cameron designed.²² Yet *Yavalath* is for sale in board game stores and is moderately successful. (As far as I know, Cameron himself pockets the profits from the game; the program that designed the game gets none of the money.) I have

myself worked on how to automatically generate video games, trying to use evolutionary algorithms. One of the systems I built tries to generate games that are learnable by other algorithms, and another system I had a hand in tries to generate games that maximize the difference in outcome between weak and strong players.²³ Both cases attempt to capture the idea that a good game provides plenty of room for different strategies and allows you to learn it. It is safe to say that none of the games that my AI systems have created have become critical or commercial successes. In particular, my team and I have frequently found that our systems will generate video games that technically should be “fun”—they are learnable and have skill differentiation—but they just do not make any sense to human players. But maybe the AI system had fun—if it could indeed experience anything, which we have no reason to believe.

So how would you combine content generation (or even game generation) with reinforcement learning to create open-ended learning? Researchers have made various attempts. Perhaps the most ambitious is Google DeepMind’s AdA (adaptive agent) system. AdA randomly generates simple games in a simulated physical environment, which are played by an agent driven by a transformer network.²⁴ The agent gets visual inputs—it “sees” the game as you do when you play a first-person shooter—and a description of the rules of the game. The game rules are very

simple; for example, the agent may get rewarded for collecting purple objects while avoiding yellow objects. Part of the reason for this simplicity is that the training process is extremely complex and therefore expensive. Still, the end result was a deep learning–based agent that is perhaps more general in its abilities than any other game-playing agent so far.

Assuming we get game generation to work better and combine it with content generation of various kinds, you can imagine how it could be used to make reinforcement learning and evolutionary computation learn better. We could keep an ever-growing list of games, and continually generate new games and game variants to challenge our best agents as they learn to master the existing games. If we gradually scale learning agents up from playing a few levels of a simple 1980s-style arcade game to being able to play a wide variety of different video games, from role-playing games to flight simulators to shooters to strategy games, it would be hard to argue that such an agent had not learned some form of general intelligence. Would a system that could play any game on the current top 100 list of iOS or Steam not be some form of AGI?

It is interesting to compare open-ended learning in gamelike environments with the large language model approach from the previous chapter and reason about what *kind* of intelligence is learned. LLMs operate in the space of text: they take text as input and produce text as output.

They don't take actions in the world and learn from them; instead they try to produce the same type of text that a human would produce given the same input. Even if they are multimodal models that can use and produce images, these images are not really connected to actions. Agents that learn in gamelike open-ended environments are very different. They operate in the space of visual observations and physical actions as they see renderings of their virtual world and act on it by moving around, collecting things, hitting things, and so on. These agents are not trying to emulate what humans do; instead they are trying to maximize their rewards, which are themselves defined by an open-ended process—not by us. Instead of learning to emulate human text production, these agents would learn to act in a world that may or may not have much in common with ours.

The potential paths to AGI discussed in this chapter and the last can also highlight some of the diverging opinions, and outright confusion, over what we mean when we talk about AGI. If we define AGI as a software system that has significantly more general capabilities for solving tasks that would seem to require intelligence than any system we have now, it is conceivable that both paths (self-supervised learning on human data and open-ended learning in virtual environments) could lead to AGI. But they would be very different kinds of AGI. One would be more of a mirror of us, operating in the sphere of our

It is conceivable that both paths (self-supervised learning on human data and open-ended learning in virtual environments) could lead to AGI. But they would be very different kinds of AGI.

cultural production (texts and images that look like the ones we produce); the other would be more of an alien intelligence, operating in virtual world(s) where it takes actions and receives rewards. While it seems as if the foundation model approach leads to more immediately useful systems, it is not at all clear which approach would lead to more general intelligence.

AGI AND CONSCIOUSNESS

Usually, when we encounter someone who shows signs of advanced intelligence, we consider them to be conscious—especially if we can talk to them. We think that there is something it is like to be that person, that they have thoughts and feelings and experiences. In fact, if you are pressed on the matter, this is how you might justify being ethical to that person. After all, if you happen to rip a T-shirt, you don't feel sorry for the T-shirt. The T-shirt doesn't have feelings. You might feel sorry for yourself because you just lost a good T-shirt. And when you explain to a child why he or she should not rip someone's T-shirt, you motivate this not with the T-shirt's feeling but with the feelings of its owner.

Now, if we developed AGI, would it be more like a person or more like a T-shirt? Would it be conscious? Would

an artificially generally intelligent system have feelings and experiences? And would we have moral obligations toward it?

The question of whether machines could ever be truly conscious has been discussed for a long time, at least since the idea of artificial intelligence was first touted. The discussion has even deeper roots than that, though. At least since Greek antiquity, we have discussed the "problem of other minds": how can we know that people we encounter are conscious?¹ It also has obvious parallels with the question of which other living beings can be said to be conscious in some way. Are other animals conscious, and if so, which? Are plants conscious? Why? Why not?

The word *consciousness* can mean many different things. The philosopher David Chalmers distinguishes between the hard and the easy problems of consciousness. Let's start with the easy problems. You might be relieved to hear that some of the problems of consciousness are easy, so we could make swift progress on them. Unfortunately, they're not really easy, more like tractable. The easy problems are such ideas as how we categorize new sensory impressions, how we can combine information so as to produce verbal report, and how impressions and memories over time can be combined to form a unified self. These are problems where we at least know what we are looking for, and we can see, at least in principle, that they should be solvable.²

The hard problem, on the other hand, is when and how these mechanisms give rise to phenomenal experience: to use Thomas Nagel's classic expression, how we know that there is something it is like to be someone.³

To see the difference between the hard problem and the easy problems, consider the possibility of a philosophical zombie. Imagine a physical body that looked and behaved just like a normal person. When asked about its past, the zombie would answer in a completely reasonable and believable manner. If you x-rayed the zombie, the pictures would look completely normal. In short, there would be nothing whatsoever to indicate that this was a zombie, not a normal person with consciousness. Yet the zombie would have no phenomenal consciousness. It would experience nothing. There would be nothing it was like to be the zombie. But if you asked the zombie if it experienced anything, it would say, "Yes, of course I do." The easy problems of consciousness relate to how the zombie can go around and talk about itself experiencing things, but also to how you and I can talk about ourselves. The hard problem is how to explain why I am not a philosophical zombie (you already know you are not one).

This discussion can appear a little dizzying or whimsical the first time you encounter it. It seems the English language is not really equipped to talk about the hard problem of consciousness. Some philosophers, such as Daniel Dennett, believe that there is no hard problem and

that the concept of a philosophical zombie is incoherent.⁴ But the idea that there is a specific philosophical problem of when and how a physical body gives rise to phenomenal consciousness is shared by many thinkers across the ages, and not only academic philosophers. The Christian concept of the soul is arguably at least partly about phenomenal consciousness.

Two of the leading theories about how phenomenal consciousness arises from matter are physicalism and functionalism. Physicalism holds that consciousness is caused by, or is a property of, physical matter.⁵ Functionalism instead says that consciousness comes from functional dependencies between the input and the output of a system or, in other words, information processing.⁶ If some version of functionalism is true, then a computer that is functionally equivalent to a biological human would be conscious if the human was conscious. Physicalism, on the other hand, says that the physical matter is crucial. Some physicalists argue that consciousness can only come from biological neurons, so it does not matter what kind of tricks a computer can perform; it can still never be conscious.

How can we tell whether functionalism or physicalism or something else is correct? Many thinkers have proposed many arguments for each of these positions. Chalmers has made several arguments for functionalism. For example, his "fading qualia" argument goes as follows: Suppose you are getting your brain replaced by a set of tiny computer

It seems the English language is not really equipped to talk about the hard problem of consciousness.

chips.⁷ One by one, small parts of your brain are removed and replaced with the chips. The chips function entirely identically to the pieces of brain they are replacing. You don't notice the replacement at all. During this process, you are awake and having all kinds of experiences; Chalmers suggests that you are perhaps watching a basketball game, being excited for the team you're rooting for and feeling the enticing smells of fast food. If functionalism is correct, then your consciousness is unbroken throughout this process, and you have phenomenal consciousness both at the beginning and at the end. If some kind of physicalism that allows only biological tissue to cause experiences is true, then your consciousness will have disappeared by the end of the procedure. But this is hard to imagine. How would this have happened—would some specific tiny piece of the brain be responsible for your consciousness, which suddenly goes away as it is removed? Unlikely. Or would your consciousness gradually fade away as, bit by bit, your brain is replaced? It is hard for us to image that consciousness can somehow fade. Can you be 0.17 conscious? If so, what would your experiences be like? Could you experience the bright red uniforms of your home team, or would they be more like a vague pink?

On the other hand, we have the arguments for physicalism. Perhaps the simplest argument, and one of the more powerful ones, is that in the history of science

everything else has been explainable in terms of physics. Quite a few phenomena, such as magnetism, the motion of the planets, and epilepsy, were at some point assumed to have nonphysical explanations, such as actions of the gods or possession by spirits. But it turns out they can all be explained in physical terms. Why would consciousness be any different?

You may not be convinced by either of these arguments and instead want to know if there is not some empirical test we can do. The problem here is that no such test exists, and probably no such test could ever exist. The nature of the hard problem of consciousness is such that it is not amenable to simple empirical tests. What would count as evidence for it being like something to be someone (or something)? The debate is very much ongoing about the status of consciousness and its basis in the physical world or functional organization, and we have nothing like a philosophical consensus. The disagreements about consciousness are every bit as vivid as the disagreements about intelligence, if not more so.

Some authors, including Chalmers, have attempted to outline criteria for when an AI system could be considered conscious.⁸ Such criteria include unified agency, memory that is shared across modalities and persists over time, and generative models of the physical world (or whatever simulated world the agent acts in).⁹ There are no current

AI models (including LLMs) that would fulfil these criteria. However, it is not clear that having phenomenal consciousness follows from having fulfilled these criteria. It is still possible to imagine a “zombie AI” that behaves as if it were conscious but isn’t really.

Where does this leave us when it comes to deciding whether an AI system is conscious? Hard to tell. Simply asking the system will not do any good. Remember that LLMs are very good at role-playing. Depending on your prompt, you can get the very same LLM to say that it is conscious or that it is not conscious. In both cases, it will answer follow-up questions reasonably. This is easy for it to do, because its training data contain plenty of discussions of consciousness and AI from philosophy and science fiction. In a sense, it is playing our own speculations about conscious AI back to us. Remember that the words produced by the LLM do not per se refer to anything, except their relation to other words. Engaging an LLM in this kind of dialogue can be entertaining, but unfortunately it does not help us understand whether it is, in fact, conscious.

It is possible that our society will at some point decide that some AI systems, or some class of AI systems, are conscious and deserve some form of rights. There will be many complicated questions about what exactly consciousness is, when, and how: The parameters of a deep network? Or its activations? Does it need to retain state over time?

Depending on your prompt, you can get the very same LLM to say that it is conscious or that it is not conscious. In both cases, it will answer follow-up questions reasonably.

What happens if you copy it? Other questions will pertain to what rights it would have: Could it be turned off? Could it be copied? Could its state be perturbed? Fundamentally, however, the decision to recognize consciousness will almost certainly be based on evidence that some people would say is insufficient. We simply do not know, and I believe we can never know.

SUPERINTELLIGENCE AND THE INTELLIGENCE EXPLOSION

In the 1984 sci-fi action movie *The Terminator*, the world has been taken over by Skynet, an AI system originally built by the US military to operate its nuclear missile defense system. Unfortunately Skynet turns out to be too intelligent and develops self-consciousness. It reasons that humanity poses a threat to it and proceeds to try to exterminate all humans, using the nuclear missiles it handily has at its disposal.

Another famous sci-fi movie, *The Matrix* from 1999, presents a somewhat similar scenario. Humanity develops artificial intelligence, but the AI system soon becomes too smart and decides it wants to run the world on its own without the interference of pesky humans. After winning the war on humanity, the AI places the remaining humans in vats and connects them to a computer simulation of the

real world circa 1999, so that the humans never know they are actually slaves to the machine.

In fact, the idea of the AI that becomes smarter than its creators and takes over from the humans who created it is surprisingly common in science fiction. In *2001: A Space Odyssey*, the 1964 movie by Stanley Kubrick and novel by Arthur C. Clarke, this happens in a spaceship on a voyage to Jupiter to investigate a mysterious transmission. More recent examples of sci-fi movies with an AI takeover theme include *I, Robot*, *Transformers*, and *Avengers: Age of Ultron*.

But this is all science fiction, which is after all just fiction. Fairy tales and fantasy contain dragons, sorcerers, demons, and deities, and just writing about them does not make any such entities exist in real life. Could an artificial intelligence that becomes smarter than us, so smart that we lose control over it and it instead starts to control us, exist? Some think so. The first person to make this argument explicit was probably the mathematician I. J. Good, in 1965:

Let an ultraintelligent machine be defined as a machine that can far surpass all the intellectual activities of any man however clever. Since the design of machines is one of these intellectual activities, an ultraintelligent machine could design even better machines; there would unquestionably be an "intelligence explosion," and the intelligence of man would be left far behind. Thus the first

ultraintelligent machine is the last invention that man need ever make, provided that the machine is docile enough to tell us how to keep it under control.¹

In other words, once we build an AI system that is as clever as we are, that system will be able to improve itself (or construct better AI systems) as well as we can. And once it has improved itself, it will be better than we are at improving itself than we are, so it will get better even faster. The rate of improvement increases exponentially, and soon the artificially intelligent machine has left us in the dust. Once this machine is literally thousands of times smarter than we are, there's no telling what it will do or how it will treat us. We might be to it as ants are to us: inconsequential and somewhat bothersome, clearly not possessing any intelligence worth caring about. So the AI system might just metaphorically walk all over us.

Quite a few people believe that this could actually happen. The philosopher Nick Bostrom wrote a book called *Superintelligence*, where he develops Good's argument in more detail and explores strategies for mitigating the risk.² David Chalmers is perhaps the most high-profile philosopher to take the argument seriously and has written an excellent analysis of it.³ Elon Musk, the head of Tesla and SpaceX, famously worries so much about the intelligence explosion that he bankrolled the (initially) nonprofit AI company OpenAI to do research that would

help us understand the risks of superintelligence and how to mitigate them. Other luminaries who have at various points taken the danger of superintelligence seriously include Microsoft cofounder Bill Gates and the legendary physicist Stephen Hawking (although he later moderated his stance).⁴ As of the time of writing, several dozen influential AI researchers have signed a short statement saying, in its entirety, that “mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks, such as pandemics and nuclear war.”⁵

Is this true? Could some kind of AI system get so intelligent that it starts improving itself, outsmarts us all, and then exterminates us, either because it dislikes us or because it just doesn’t care and could use the atoms we are made of for some better purpose? Let’s consider these two questions separately: Can AI software self-improve radically, and can it pose an existential risk to humanity?

In some sense, it is easy to make a self-improving AI system. Self-improving AI has been around for decades. For example, one can implement an evolutionary algorithm that also evolves its own hyper-parameters (its “settings”) to be more efficient. Or you can use reinforcement learning to make a neural network produce better neural network structures for producing better versions of itself. However, this type of self-improvement is not what we’re looking for here. These systems have a low ceiling, meaning that their self-improvement only takes them so

far. They are good at optimizing some aspect of their own functioning to a degree, but cannot add completely new functionality. (Part of the reason for this is that there is no way of training a system to invent completely novel functionality; by definition, there is nothing to train it on.) What the intelligence explosion argument argues for is systems that can get drastically better very fast, so that we lose control—basically, systems that could, on their own, jump years ahead in AI development.

The intelligence explosion argument relies on the concept of intelligence. Because the system is intelligent, and intelligence is what we need to create even better AI systems, the system can use its own intelligence to become more intelligent. But, as we have seen, what we mean by intelligence is fuzzy at best. Will the same type of intelligence that the AI system has be useful for making the AI system more intelligent? And will the same type of intelligence that is needed for upgrading the AI system one step be the type that is needed for upgrading it to the next step?

Actual AI development as practiced by humans requires a whole range of different skills. The most visible ones are in designing the architectures and training methods, for example, neural network structures and optimization algorithms. But these skills are only the tip of the iceberg. Every bit as important is building the data and computer infrastructure so that these extremely complex systems can function adequately. Then there are collecting and curating

the training data. And, of course, building the hardware. As mentioned earlier in the book, advances in hardware (including CPUs, GPUs, memory, storage, and networking) have been crucial to enabling modern AI. The latest advancements could not have been made with last-generation hardware, because training would have been too slow, or memory would have been insufficient. So for an AI system to improve itself beyond a small increment of its current capabilities, it would need to be able to make progress on multiple fronts. This is a lot to ask even on the software side, and it would be complicated, to say the least, on the hardware side. At the least, it could not be done quickly. Manufacturing new chips (as in processors or memory) is an extremely complex business with a long international supply chain where only a few companies in the world can produce each component. It takes months in the best case, and an AI system trying to do this on its own would be very obvious. Accordingly, a rapid and sustained capability increase for an AI system *on its own* seems highly unlikely.

Even more importantly, progress in developing new AI is not just a matter of intelligence. Like all forms of research and development, it requires experimentation. We never know what is going to work in advance. We need to try things and fail, and trying things takes time; the most advanced recent large language models took months to train, and even longer to experimentally find the right architecture and settings that allowed training to work.

A rapid and sustained capability increase for an AI system *on its own* seems highly unlikely.

In sum, it is more than imaginable that an AI system can improve its own capacity to some extent. Rudimentary versions of such systems have existed for a long time. But for an AI system to improve its own capabilities by more than a small amount would require an extremely broad knowledge base, access to near-unique manufacturing capacity in multiple countries, ability to conduct empirical experiments, and lots of time. None of that happens by itself, or fast and undetected.

Of course, we could imagine an AI system self-improving with the help of a large team of smart and dedicated humans who have the right expertise, connections, and financial resources. But that is essentially how AI software and hardware are developed right now. They are developed by large teams of humans, in various organizations, with the help of AI tools. But it makes more sense to consider the humans, rather than the AI systems, as the agent. Or, perhaps more appropriately, companies, or even our technical civilization in its entirety, is the real self-improving agent.

Let us move on to the next question: even if we don't think that AI can self-improve by more than a marginal degree, can an advanced AI system pose an existential risk to humanity?

The scenario from *The Terminator* is not impossible. (At least the first part, where it unleashes a nuclear war; the part with the time-traveling humanoid robots is more fanciful.)

But such a scenario requires that the system is agentic, and building agentic systems is a choice. More importantly, it requires that someone has given an agentic, uncontrollable AI system complete control over the nuclear launch sequence. This is obviously not a great idea. But the problem is more of a security problem than an AGI problem. All of this could happen *without* AI. Since nuclear weapons were invented, we have had several close brushes with accidentally triggering a nuclear war because of faulty procedures. The 1984 movie *War Games*, one of the best movies about AI ever made, sets out a scenario where a less-than-general AI almost triggers a nuclear war. It is believable because the AI system in the movie is actually not that sophisticated. The lesson here is less about AI, and more about having proper security procedures in place with humans in the loop.

Bostrom lays out a scenario where an AI system destroys humanity without wanting to, or even necessarily having anything against humanity at all. In Bostrom's scenario, an AI system is given control of something as innocuous as a factory that makes paper clips.⁶ The AI system has a single goal: to make as many paper clips as it can. The system then becomes superintelligent and proceeds to invent new ways of maximizing its paper clip output. It gains control of all humanity and converts all our manufacturing output and eventually us as well into paper clips.

This story hinges on the plot device that the AI suddenly "wakes up" and becomes much more capable than it

Building agentic systems
is a choice.

was before, and too capable for us to control. But, as I argued earlier, this seems unlikely at best. It is also unlikely that even if you somehow had control over every machine in the world that was connected to the internet, you could use that to turn us all into paper clips.

You could still imagine an extremely capable AI system causing real harm to society even without a sudden capability increase, if you assumed that it was the only AI system of that magnitude of capability. This is a cybersecurity problem more than an AI problem; if you imagined that a hacker of today had access to all existing hacking tools, but nobody had any cybersecurity defense, that hacker could take control of plenty of systems and cause havoc. Fortunately, however, cybersecurity defense and offense capabilities have historically evolved gradually, because this knowledge is spread among many actors in society. It seems that the safest way to guard against AI-enabled cyberattacks is to make sure that modern AI capabilities are spread among many actors in society.

This chapter has discussed the more imaginative ways that AI systems could mess with us. As you can tell, I am not particularly worried myself, although some others are. This is a discussion we will keep having for a while. The next chapter, in contrast, will discuss the more pedestrian, or at least less catastrophic, ways in which more general AI may affect society.

AGI AND SOCIETY

It is hard to see how the availability of more generally capable AI could not have large impacts on society. Which impacts specifically depend on what kind of AGI we are talking about, but of course also which society and how it handles the disruptions caused by the new technology.

The discussion in this chapter will by necessity be even more cursory, shallow even, than in the rest of the book. This is partly because of the complexity of the topic, and partly because of my level of expertise (or lack thereof) in the subject matter. I am not a social scientist, just a lowly computer scientist with eclectic interests and a background in philosophy. Like the rest of the book, this chapter takes the view that AGI is simply more competent and general AI technology than we have now, so much of the discussion here concerns more narrow AI as well.

The first concern that many people have when discussing the societal impacts of AGI is job displacement. In modern societies, people work to make money, which they then spend on goods and services produced by others. If an AGI system could do everything that a human can, what is there for us to do? How will we make money? And how will we fill our lives with meaning, given that work is important to many of us for more than pecuniary reasons?¹

A world where AGI does all the work and humans do other things with their lives has been explored in various science fiction novels. Perhaps the most complete exploration is in Iain M. Banks's *Culture* series, which comprises nine novels centering on characters connected to "the Culture," a powerful society where humans live in symbiosis with highly advanced technology.² The crown jewel of this technology is the Minds, extremely powerful AGI devices—or beings—that are a million times as intelligent as a human and can do any intellectual task of humans better. This is the archetype for the "disembodied mind" type of AGI discussed in chapter 5. Helping the Minds are a vast variety of mobile robots, making sure that all material needs of humans are taken care of, including medical care; humans essentially live for however long they want. For most humans, life in this society is good. Humans do whatever they want and have no particular expectations on them, except not to hurt each other. Generally humans engage in friendship, food, romance, sex, and

various games, hobbies, and pastime activities. They may spend a few decades going deep into some philosophy, art, or craft, knowing that they can always ask the Minds for help, because the Minds always know better. Thrill seekers engage in extreme sports or travel the universe. Some join Contact, the foreign service of the Culture, to engage in diplomatic or nondiplomatic contact with other species and civilizations.

The Culture novels are well worth reading. They have plenty of interesting speculation, some good plotlines, and good action too. Still, reading them makes you wonder whether actual humans would feel content with such an existence. If we had no strife and no challenges to overcome, would life not feel a little empty? Most of the events in the novels take place in the interface between other civilizations and the Culture. That's where the tension, and the action, is. It is apparently hard to write good stories that take place in an AI-driven paradise.

Back in the world where we actually live today, it is clear that to the extent that AI systems will replace humans at their jobs, it will be a gradual phenomenon. It will also be highly uneven, with some jobs and tasks more exposed than others. Interestingly, we seem to be witnessing an inversion of a trend that has been holding since the Industrial Revolution, where manual labor gradually got automated and people moved on to desk jobs. Recent advances in AI largely concern understanding and generating text,

code, voice, and images, exactly the kind of things that desk workers (or these days, remote workers) do. In contrast, complex manipulation of physical objects is largely beyond the state of current robotics. We do not have reliable self-driving cars on our highway, much less useful humanoid robots. No robot today can replace, or even assist, a plumber. The industrial robots that perform tasks in factories are highly scripted and the opposite of general intelligence. This phenomenon, where AI technology performs better at the kind of tasks that we would consider cognitive and much worse at supposedly nonintellectual manual tasks, is called Moravec's paradox and has been acknowledged for decades.³

Throughout recorded human history, humanity has continuously made technical and organizational innovations, and these innovations have transformed or in some cases replaced jobs. The Industrial Revolution was partly a consequence of technological improvement in farm labor, meaning that large numbers of workers were available for working in the new factories in cities. Work in the factories was gradually automated, and workers shifted to overseeing the machines, and then to desk jobs. When we could automate or partly automate a job, new jobs were created as new demands could be met. Historically, unemployment has temporarily shot up during times of great technological change but then receded as workers have transitioned to new roles. When you average over these

To the extent that AI systems will replace humans at their jobs, it will be a gradual phenomenon. It will also be highly uneven, with some jobs and tasks more exposed than others.

technological cycles, unemployment has been surprisingly stable for hundreds of years. The question is why things would be different this time, especially as AI technology as we know it is a gradual evolution of computer technology that has been in widespread use for decades.

To take some recent examples of jobs that were not replaced, we have bank tellers and accountants. When ATMs (automated teller machines) were introduced, it was predicted that banks would fire a large number of their human staff. Instead the number of bank tellers grew, because more people now had access to and demand for bank services.⁴ Similarly, some predicted that accountants would become obsolete as spreadsheet software (such as VisiCalc, and later Microsoft Excel) was invented, but what ended up happening was that accountants could do more with less and were in more demand than ever.⁵

In most cases, it is more fruitful to think of human jobs being augmented by AI technology than replaced by it. In the same way that writers and journalists have had to learn to use, and eventually become dependent on, word processors and spell-checkers, they will probably learn to use LLMs for writing assistance. The increased productivity this affords will become expected of them, and AI-assisted writing will become the norm. The same argument can be made for many professions: advertising professionals and editors had to adapt to desktop publishing technology in the 1990s, which streamlined the process

of producing printed pages but did not reduce the number of employees in the advertising and publishing industries. It's true that old-timers who stuck with manual production processes did gradually lose their jobs, but those who learned desktop publishing became more productive than ever before. At the same time, the lowered cost of producing printed media (and web pages, as the same content was increasingly published both on paper and online) meant that the demand for publishing and advertising became higher than ever.

The same argument can be made for blue-collar professions as well. For example, professional drivers have been widely discussed.⁶ Truck and bus drivers are some of the most common workers in many countries, including the United States. These jobs have already been through many cycles of technological advancements as trucks and buses have become larger and more reliable and payment and routing solutions more automated, meaning that the same amount of cargo or people can be moved with fewer drivers and other people involved. Now there are worries that AI technologies will simply make drivers superfluous. Although the technical hurdles seem larger now than they did ten years ago, it is likely that we will eventually have self-driving trucks and buses in our cities. What will drivers do then? Well, these self-driving buses and trucks will still need to be managed by someone, in the sense that humans will make all the business decisions and do all the

paperwork. Trucks will probably need help loading and unloading. What's more, someone will need to monitor these fleets of self-driving vehicles (probably remotely) and intervene when problems occur, as they always will. And let's not forget customer service, maintenance, software development, and so on. And as it gets easier, faster, and cheaper to transport things, it is likely that demand for transports will go up as well. In short, it is likely that self-driving vehicles will create lots of jobs at the same time as they massively increase our capacity for moving goods and people. Many of these jobs will require knowledge of vehicles, traffic, and the peculiarities of individual cities and regions, which is exactly the kind of knowledge that professional drivers have.

We haven't even discussed all the new jobs that AGI could enable. There is a near-limitless number of tasks that are not done because we cannot afford to do them, or because we are busy doing other things. I am a little bit disappointed that our apartment does not clean itself, that I still have to book flights and sync meetings myself, and that my wife and I have to do the annoying parts of childcare (making sure our toddler son doesn't do anything dangerous, such as running down the stairs) instead of just focusing on the fun parts. I'm also annoyed that I have to do my taxes, that I can't get a professional-quality massage whenever and wherever I want, that I don't have a high-quality medical examination at hand immediately

if I get bitten by a tick or wonder about that weird feeling in my left knee, that I don't have a master shoemaker fixing my shoe right away if it's slightly uncomfortable, that I don't have a professional bartender mixing me a cocktail right now, and so on. In case I would like to rap, it would be nice to have someone like Snoop Dogg there to coach me. These are things that are available to very rich people (through paying humans) or the denizens of Banks's fictional Culture (through AGI). The point is that human demand is practically infinite, and combinations of AI technology and skilled human labor are going to make many more goods and services available to many more people than ever before. It seems highly unlikely that we would run out of things that we would want to pay people for.

In sum, history strongly suggests that we should not worry about long-term job losses due to more generally capable AI. I would argue that this is true even if we eventually get the kind of AGI that is better than we are at literally everything. Our jobs would then be to be the bosses of this marvelous technology, because in the end we are the ones who best know what we want. This is consistent with history, where technology is increasingly doing the tasks we used to do, and we are managing the technology. Even so-called low-skilled jobs now often consist in pressing buttons, telling machines what to do. We, as a civilization and as individual laborers, are constantly moving up

There is a near-limitless number of tasks that are not done because we cannot afford to do them, or because we are busy doing other things.

the value chain from repetitive and simple manual labor toward decision-making.

In an interview on CBS's flagship show *60 Minutes*, the star journalist Anderson Cooper interviews the legendary music producer Rick Rubin. Throughout a long career, Rubin has produced a large number of artists in different genres, ranging from Beastie Boys and Jay-Z to the Red Hot Chili Peppers, Aerosmith, Rage against the Machine, and Johnny Cash. Rubin's ability to help these artists reach higher levels of their potential is beyond doubt. He's a legend for a reason. In the interview, Cooper asks Rubin if he plays any instruments, and Rubin answers "barely." Rubin goes on to state that he doesn't know how to use modern recording equipment, has no technical ability, and knows nothing about music. Cooper then asks Rubin what he's being paid for. It is worth quoting Rubin's answer in its entirety: "The confidence that I have in my taste and my ability to express what I feel has proven helpful for artists."⁷

I think about this interview often. It seems to me that Rubin's job is remarkably AI safe. Of course, you could use some sort of AI system to give feedback on a piece of music and suggest changes. But that would still not exactly be Rubin's taste. And that's the point. In the event that we ever get AGI systems that are able to do everything we do better than us, maybe we will all have Rubin's job, or the equivalent for whatever domain we work in.

I have dwelled on labor and potential job losses because work fills such a large part of our lives, but more general AI could affect society in countless other ways, far more than I could do justice to in this short space. It is worth briefly visiting three other topics: entrenched biases, misinformation, and cybersecurity.

The problem with entrenched biases is that, to the extent AI systems are trained on data generated by humans, AI will learn to reproduce the biases inherent in those data sets. By now, plenty of evidence shows that this is happening. One of my favorite demonstrations of this phenomenon concerns the widely used word-embedding capabilities of Word2vec. Word2vec is a neural network trained to map words into a multidimensional space, which is useful in a wide variety of natural language processing applications.⁸ One of the nifty things about such word embeddings is that you can do a kind of arithmetic using them. For example, you can take the embedding for *king* and subtract the embedding for *man*, and remarkably the embedding for *queen* emerges. The problem with this is that, for example, if you take the embedding for *doctor* and subtract the embedding for *man*, you get the embedding for *nurse*.⁹ Note that these biases have not been consciously programmed by anyone but simply learned from the data, which are a snapshot of text from the web. Nevertheless, a system that makes such assumptions about gender is going to say a number of things, or

perhaps make decisions, that are far from what we want it to do.

This phenomenon is not constrained to simple language models but is pervasive in large foundation models as well. You can try for yourself to ask a sophisticated image-generating model such as Stable Diffusion to generate images of doctors or nurses, or for that matter scientists or plumbers, and see what gender and ethnicity the people in the generated images are.¹⁰

One possible reaction is that the trained models simply reproduce the structures of our society, and we should let them do that. But this overlooks that we often want the models to specifically *not* do this. For example, we may want our illustrative images to reflect a diverse and inclusive society, not the power structures of yore. Or we may want the models that assist with our decisions on mortgage approvals or hiring to specifically *not* take gender, ethnicity, or age into account.¹¹ This is often also legally mandated. Of course, you could argue that we should not rely on AI systems to make such important decisions at all, but given the general applicability of this technology, chances are that we will. Thinking is hard, and most of us conserve our limited thinking resources when we can.

How can we make sure that models are less biased? Perhaps the most obvious route is to not train them on biased data. But this is hard, given that most types of models benefit from being trained on as many data as possible,

and so many existing data are biased. Bias may also be encoded in subtle ways, such as how often women are mentioned compared to men in certain contexts, meaning that discarding biased data might require complicated and labor-intensive analysis. There are also many different types of biases (beyond ethnicity and gender, there are age, looks, religion, politics, and many other categories) that are expressed in many different ways, meaning that accurately filtering training data might mean doing lots of work to end up with fewer data and worse-performing models.

Another approach is to try to adjust biases after the model has been trained. RLHF (reinforcement learning from human feedback) has been used with some success to debias large language models, but unfortunately it seems that this process makes the models perform worse in other ways.¹² In either case, efforts to limit, mitigate, and even detect bias are an active research field.¹³ It is safe to say that we have much progress left to make here, but also that it is highly unlikely that we will find a simple solution that does away with unwanted biases without causing other problems.

The impressive capabilities of modern foundation models to generate high-quality text, voice, and images come with the alarming possibility for generating high-quality misinformation.¹⁴ For example, it is easier than ever to generate photos of an existing person engaging in

something that never happened, whether a crime, a sexual act, a battle, or even just a party with the wrong kind of people. This is a problem because we are accustomed to believing photos. It is true that people have fabricated photo evidence almost since cameras were invented, but it used to be hard work with results of varying quality. Think of the now-infamous photos of early Soviet leadership where Leon Trotsky was removed after he was expelled from the government. Those quite inept edits were the best a major world power could produce once upon a time. It's now easy to create convincing fake photos, and the technology is advancing rapidly. (In a previous chapter, we discussed the picture of Pope Francis wearing a stylish jacket from the fashion brand Balenciaga generated with Midjourney; the picture went viral, and new versions of image-generation models have since been tested by generating blinged-out Balenciaga-wearing Popes.)¹⁵ It is very likely that we will soon also be able to generate convincing fake videos. Fake text poses a different kind of challenge: while we are not typically likely to believe something just because a random person on the internet said so, we may be convinced by what looks like popular consensus. LLMs make it easy and cheap to create thousands of posts propagating a political opinion, claiming that an event has happened, or praising or criticizing a person or product. To the extent we rely on the opinions of anonymous or pseudonymous people online, this poses a problem for democracy.

Most likely, society will adapt to cheap and plentiful misinformation just as it has adapted to other information disruptions (the printing press, telephones, TV, the web). An urgent question is how we can help society adapt rapidly while minimizing disruptions. After all, you could argue that the printing press, with its ability to print cheap pamphlets, enabled the Lutheran Protestant Reformation, thereby indirectly causing the Thirty Years' War, which killed an estimated 20 percent of Europe's population. Note that I am *not* arguing that the printing press should not have been invented or been allowed to proliferate.

Finally, a few words on cybersecurity. LLMs in particular have shown impressive capabilities in writing code and even in adjusting the code they have written based on the results of running it. While this is useful in many circumstances, it would also seem to make it possible for LLM-based agents to semiautonomously hack into other systems. This is worrisome, as it could enable malicious actors to unleash rogue agents that do considerable damage. A related source of worries is that more powerful and general AI methods can be used to attack our systems for authentication and identification, including passwords but also biometrics based on analyzing your face, voice, fingerprints, or iris, for example. Some systems already use modern machine learning to generate fake biometrics that can fool many biometrics systems into authenticating a hacker as a legitimate user.¹⁶

In cybersecurity, it is generally accepted that threats evolve and defenses need to evolve with them. We need to use a plurality of defensive measures to protect valuable or vulnerable systems. Cybersecurity is based on constant evolution of attacks and defenses, and AI methods are increasingly used for both. In other words, the threat of AI-based attacks is best handled by AI-based defenses, meaning that the dispersion of AI systems and capabilities widely in society is very much in our interest.

CONCLUSION

Writing a book about artificial general intelligence in 2023 means chasing a moving target. I started working on this book before ChatGPT was released, quickly gained tens of millions of users, and reinvigorated the AGI debate. As I am finalizing this book in October 2023, there is intense discussion among both technologists and politicians about AI regulation. By the time you read this, much will have happened in terms of both technology and policy. But my descriptions of how the underlying computational methods work will still be true, and I think the various points I make throughout the book about both natural and artificial intelligence will still apply.

What are those points, exactly? In this final chapter, I attempt to synthesize the themes of the book and to look forward. I also allow myself to express some of my opinions (I have many) about AI in general and general AI in

particular. Let us start from the beginning and summarize and comment on what we discussed in each chapter.

Artificial intelligence is not a single thing. The term refers to a number of different approaches to building “intelligent” software, meaning that it solves problems that seem to require intelligence. Many approaches to building such software exist, including various forms of planning and tree search and machine learning, for example through neural networks. The best approach depends on the problem. Since usable digital computers were developed in the 1950s, we have seen a number of AI systems that showed superhuman capabilities on some particular problem. Early on, we had mathematical theorem proving, later image recognition, and then chess and various video games. More recently we have systems that can analyze and generate images and text, maybe not better than a human professional but certainly faster. These systems are all in some ways specialized; the best language model is not very good at proving theorems and playing chess, and the best chess-playing system would not even know how to begin to write a poem. None of these systems could tie their shoelaces if they had any.

What is so elusive is thus artificial general intelligence. It is so elusive that the definition of the term is itself elusive! It turns out we don’t have a good definition of general intelligence or even of intelligence. The definitions used in psychology are at their core based on the results of various

Artificial intelligence is not a single thing. The term refers to a number of different approaches to building “intelligent” software.

cognitive tests. IQ simply measures how good you are at IQ tests, though it is also pretty good at predicting a number of other things, such as career advancement in many professions. But which tests exactly should be included when characterizing intelligence is to some extent a matter of opinion. Looking beyond human psychology, we see that other animals exhibit intelligences of vastly varying character. We have no obvious way to rank a dog, a parrot, and an octopus on the same intelligence ladder. This makes sense, as these animals evolved in different ecological niches, and their brains and bodies are adapted to the specific challenges of that niche. While it may seem that it should be easier to come up with a definition of general intelligence from mathematics and computer science, it isn't really—at least if you want a useful definition. Universal intelligence is a nice mathematical idea based on summing performance over all possible environments, focusing on those that are in a sense simpler. However, it is impossible to compute. It is also not clear that any existing intelligence, including that of humans, would score well on this measure. And many of the “environments” in which humans perform span large, overlapping parts of our life, with unclear rewards.

Perhaps it is not a big deal that we cannot define artificial general intelligence. Perhaps it is more interesting to look at the various visions of more-general-than-we-currently-have AI systems that have been proposed by

scientists and science fiction authors, and try to categorize them according to which traits the potential future AGI systems would have. Doing this, we see that visions for our AGI future differ markedly: will an AGI system be a single entity or a collection of interacting agents? Will it be embodied (as a robot) or not? How much learning will it need to master the same skills we have? How generally skilled will it be? Will it have a will of its own, or is it more like a collection of services, a Microsoft Office with genius-level skills on demand? And perhaps most importantly, how humanlike will it be? It is easy to fall prey to the idea that an advanced AI system would behave and think similarly to us, simply because we are the only frame of reference we have for advanced and (somewhat) general intelligence. Conversely, it is difficult to imagine a truly different intelligence. Authors who try, like Stanisław Lem or China Miéville, mainly succeed in portraying our own reactions to alien intelligences. It is not at all clear that an AGI system would be humanlike, and this would at least partly depend on how we develop such systems; if they are developed largely by imitating humans, they might be more humanlike, but if they are developed by setting and solving their own goals, they might become quite alien to us indeed.

As for how to build AGI, there is no shortage of opinion. Throughout the history of AI, a number of different methods have been invented. These can roughly be categorized

as variants of search, optimization, and learning. Which method works best depends very much on the nature of the problem, but none of these methods seem to be sufficient to build significantly more general and capable AI systems *on their own*. Combining these methods is a different thing.

As I am writing this, self-supervised learning of foundation models, such as large language models and image generators, is all the rage—and probably will be for some time to come. LLMs such as GPT-4, Llama, Claude, Gemini, and whatever the new language model du jour is when you read this are fundamentally trained to predict the next word on very, very large amounts of text. Similarly, image generators such as Stable Diffusion, Dall-E 3, and Midjourney are trained to re-create images from noisy versions of these images. The words *self-supervised* and *autoregressive*, commonly used to describe these kind of models, simply mean that the models are trained to produce the data from themselves. The surprising thing is how well this works, and how these models can do all kinds of unexpected things. Image generators can create elaborate images that have never existed before while closely adhering to complex prompts describing nonexistent situations. LLMs can do . . . a lot. It turns out that many tasks can be rephrased as text completion, a souped-up version of the autocomplete functionality that keeps mangling your text messages. Part of the trick here is instruction tuning,

where the model is additionally trained on sets of questions and answers, or instructions and results. A properly trained LLM can answer questions, solve puzzles, write love poems, and many other things with varying output quality.

Seeing the power and versatility of LLMs, some people claim that they are the first steps toward AGI. I don't agree, and not only because we don't have a good definition of AGI. LLMs have plenty of shortcomings: they are generally bad at reasoning and planning, can't (on their own) really do math, and, perhaps most importantly, are extremely unreliable. Even the best LLMs hallucinate frequently in some situations. This makes them unsuitable on their own in many situations. However, it is possible to build systems around LLMs that amplify their power in various ways. For example, one may let them write code, do internet searches, call math software, and other things. It is possible that something we might call AGI could one day emerge from systems where LLMs are components.

A completely different approach to AGI is represented by open-ended learning. Here we are not training models on the cultural exhaust of humanity (text and images). Instead we build coupled systems of agents and problem generators, loosely inspired by how natural intelligence has developed through evolution in complex ecologies. This line of AI development is partly inspired by video games and often takes place in video-game-like simulators

or even actual video games. An AI system can generate problems (such as game levels or game rules) for agents to learn, and when the agents get good enough at that problem, new problems are generated. Bit by bit the agents—or perhaps societies of agents—learn to solve more complex problems and thrive in more challenging environments. I believe that open-ended learning is a viable path to some form of AGI. However, it would almost certainly be a very different form of AGI compared to intelligence based on self-supervised learning on human-made texts and images. We have already seen examples of this when open-ended learning methods produce video games that fulfill the criteria we made for being “good,” but are largely incomprehensible to human players. This challenges the notion of general intelligence: how general? In relation to what? While humans naturally have a strong anthropocentrism, I think we have much to learn from AI systems that are fundamentally different.

If we can build an AGI system, would it be conscious? Would there be something it would be like to be that system? Would it feel genuine pain and happiness, not just perform simulations of those feelings? We don’t know. Several generations of philosophers have tried to answer this question and proposed many clever arguments, and philosophical arguments about what kinds of creatures or things have minds have been around for 2,500 years or so. I personally think this is one of the most important

questions in the world, and I majored in philosophy in my undergrad largely because I wanted to understand whether intelligent machines could or would be conscious. After a few years of thinking about this, I concluded that I had no idea how I (or anyone else) could make progress on this question, so I switched to the easier problem of trying to create artificial intelligence. And that's still where I am. I admire those who try to answer this question, however.

A question that is surprisingly, and annoyingly, widely debated as I'm writing this is whether AI risks leading to human extinction. The idea here is that once we create sufficiently advanced AI systems, they will be so intelligent that we cannot control them. They might proceed to improve themselves to be even further beyond our grasp, and then we might be to them as ants are to us. With their inscrutable priorities, they might just decide to kill us for some reason. I know, this sounds suspiciously like science fiction or even theology. The problem with the argument, and the debate around it, is that it builds on a number of ill-defined concepts. Chief among them is "intelligence" (sometimes near synonyms are used), which, as we have seen in this book, is a concept that stubbornly escapes meaningful definition. If you try to use well-defined words that refer to actual AI technologies and their capacities, the argument falls apart.

It is interesting to compare the argument that AI poses an existential risk with other existential or civilizational

risks that people worry about, such as nuclear war, runaway global warming, asteroid strikes, and pandemics. In each of these cases, we understand the mechanisms behind the potential catastrophe in some detail. We also have some historical examples to learn from, for example, the Hiroshima and Nagasaki bombs, the Permian-Eocene thermal maximum, the Tunguska event, and the Black Death. In other words, we know that these risks are real, and we know quite a few things about how they work. In comparison, the AI existential risk argument is mostly fluff. It's true that we should not discount risks just because we don't understand them well. But conversely, not understanding something does not mean we should worry that it carries an existential risk.

Unfortunately, some people are convinced that future AI development poses serious risks and should be curtailed and controlled. Others are interested in bogging down AI research and development with restrictions and governmental bureaucracy to preserve their own competitive position in AI development and therefore claim to worry about existential risk. It's not pretty. As I'm writing this, proposals are being floated to require licenses for training large models. I think such regulations would be a big mistake, not only because of the chilling effects on AI research, but also because they represent a slippery slope in terms of freedom of speech. Part of the issue is that AI, as we have seen, comes in many forms, and we don't know

what shape the next groundbreaking AI method will take. As we have seen in this book, even defining AI is very hard; the best attempts at definition tend to be either specific to a particular method (and thus have a short shelf life) or so broad as to encompass lots of things that we don't think of as AI. Think of word processors and pocket calculators. So, to be effective, any regulation of AI methods would need to be rather broad and intrusive. Now ask yourself whether you want there to be laws about which programs you can write and run on your own computer? And how, exactly, would such laws be enforced? By some mandatory government spyware? I am not saying this is likely to happen anytime soon, but the concern over existential risks from AI seems to lend itself to such proposals.

This is not to say that there are no risks associated with the development and diffusion of more capable and general AI systems. Given the wide applicability of anything we may call AGI, it is likely that it will have major effects on society. While some worry that many, or even most, people may lose their jobs, I don't think that is likely. But I do think that many, maybe most, people will need to adapt to using various forms of AI support in their jobs to stay competitive. Maybe we'll all get accustomed to using LLM-based agents to do much of our bidding. Or some other kind of technology—as we have seen, it is hard to predict which kind of method will dominate the next wave of AI.

While I am not worried about some future AGI system waking up and deciding to make paper clips out of all of us, I do think it is worth worrying a bit about the cybersecurity aspects of powerful AI systems. When almost every part of our technical civilization depends in some way on being connected to the internet, cybersecurity is an important matter. Bad actors—states, terrorists, doomsday cults, or just plain old criminals—can do a lot of damage using AI-enhanced hacking. This can occur both in the strictly technical sense, where software vulnerabilities can be automatically found and exploited, and in the social-engineering sense, where convincing fake personalities are created to trick you into giving away important information. Then again, AI methods also make new defenses possible. This is not hypothetical; it's already happening. It is probably best seen as an extension of the ongoing race between attackers and defenders that has defined cybersecurity since the beginning. Various AI methods now provide tools for both attacking and defending systems.

Our computer systems are relatively safe today because so many software developers and system administrators take security so seriously and share their knowledge freely and openly. Cybersecurity is studied in both academia and industry, and papers and source code are shared openly. You might think that this would give an advantage to attackers, but the opposite is actually true. The best way to test your defense strategy is to let other hackers or

researchers try to attack it and share what they learned. And the best way to rapidly enhance security methods is to let hackers and researchers freely build on each other's solutions. Besides, trying to stop bad actors from sharing attack strategies with each other would clearly be futile, so it makes sense for good actors to employ the same strategy.

I think this is the future we want for AI methods and models as well. As much of AI research and development as possible should be conducted openly and transparently. This means not only that the model parameters and the code used for training should be open-source but also that researchers and developers openly publish their methods and findings. This approach will allow new models and methods to be tested by anyone in the world with the technical capacity, leading not only to more innovation but also to more safety. A society where as many people as possible have access to, understand, and can contribute to modern AI will be safer from whatever risks AI systems might bring.

Finally, before I let you go, I would like to reflect on what it means that the terms *intelligence*, *artificial intelligence*, and *artificial general intelligence* are so stubbornly hard to define. Does this mean that we just haven't understood what intelligence or artificial intelligence is yet? But that assumes we are studying a natural kind when we study intelligence, that is, a grouping that reflects something real in nature. I don't think that is the case. I think

intelligence is a word we made up to represent a somewhat arbitrary set of capabilities that humans tend to possess. Trying to find the “core” or “essence” of intelligence is probably a fool’s errand, much like finding the essence of funkiness, beauty, or the 1980s. There is no secret sauce, fundamental principle, or “one simple trick” to intelligence. Things don’t get any more definite if we place the word *artificial* before *intelligence*. Artificial intelligence was just the name of a seminar in 1956 that somehow also became the name of a sprawling research field and the various technologies that emerged from it. That the various technologies discussed in this book are all referred to as AI is mostly a historical accident and/or marketing. We could have used a different term or several different terms. In fact, some of the early neural network research was done under the moniker “cybernetics.”

One consequence of all of this is that there is not really anything magical about AI. It’s just a set of useful technologies, some of which might change the world—in the way that clocks, stirrups, steam engines, and telephones all once did. This realization is a little painful for those of us who chose to become AI researchers because of that magic, those of us who wanted—and in some sense still want—to understand the mind by creating minds with computers. There are indeed plenty of interesting technologies to develop and phenomena to understand. There’s just not any great mystery to solve.

Trying to find the
“core” or “essence” of
intelligence is probably
a fool’s errand, much
like finding the essence
of funkiness, beauty, or
the 1980s.

As for AGI, I think the term is only useful to refer to a particular direction of AI research, toward more generality. Statements about if or when we will reach AGI are meaningless because of the lack of good definitions. In fact, I think it would be best if we all simply stopped talking about AGI. It is leading us astray from the more important questions, which tend to focus on particular applications of AI technology and their consequences for society.

Admitting that intelligence is a rather shallow concept and that there is nothing magic about AI allows us to see the history of AI from a new perspective. Throughout this book, I have discussed AI as a series of technical inventions motivated by being able to solve problems that require intelligence. This history has been characterized—some would say plagued—by the moving-goalpost problem, where we set up a goal to work toward, and when we have developed AI methods that reach that goal, we say that it isn't good enough, and we create a harder goal. And so the goalposts have moved from chess to Go to general video game playing, and from simple object detection to segmentation of images and identification of multiple objects in low-quality photos. As the saying goes, AI is just computer science that doesn't really work yet; and when it works, it's no longer AI.

The alternative perspective is that the history of AI is a long deconstruction of the concept of intelligence. This proceeds by someone confidently exclaiming that

something—say, planning, image creation, or translation—is a hallmark of real intelligence. The research community then finds a way to perform this feat using some new or old AI method. We then look at the original task and say that it didn't really require intelligence after all, because it can be done with mere computation. Therefore we need to find another feat that really requires intelligence. And then we do this again and again, chipping away at the concept of intelligence. The job will be done when we can no longer find anything that we can claim requires intelligence and that we cannot get a computer to do as well as we do it. The concept of intelligence will then become pointless, except in the colloquial use of the term. At that point, we may or may not want to say that we have achieved AGI.

In Alan Turing's brilliant 1950 paper "Computing Machinery and Intelligence," which proposed the Turing test and has been a cornerstone for generations of AI researchers, Turing sets out to answer the question "can a machine think?"¹ At the end of the paper, having gone through many arguments for various positions on this question, he asserts that the original question is "too meaningless to deserve discussion." In other words, carefully analyzing the question reveals that it is badly posed. I think this is also the case for the question of if and when we can achieve artificial general intelligence.

ACKNOWLEDGMENTS

I received valuable comments on a full draft of the manuscript from several kind and knowledgeable reviewers: Sam Earle, Noah Giansiracusa, Ahmed Khalifa, Jerry Swan, Mike Preuss, and three anonymous reviewers. Their suggestions have helped me improve the book in many ways, even when I didn't implement them exactly as suggested. Of course, all errors and mistakes are entirely my own responsibility. However, I am pretty sure this text is more factual than if GPT-4 had written it. And, believe it or not, at no point did I enlist the help of a language model for writing this book.

I also want to thank Elizabeth Swayze and Matt Valades at MIT Press, who supported the project enthusiastically and had patience with my liberal attitude toward deadlines.

Above all, thanks to Raz and Dylan for putting up with me spending late nights and sunny weekend afternoons with this book instead of with them. I love you!

GLOSSARY

Artificial general intelligence (AGI)

An AI system that has significantly more general capabilities than any system we have now. Or, alternatively, an AI system that is capable of performing most economically valuable tasks that humans perform.

Artificial intelligence (AI)

A research field that studies how to build technologies that can perform tasks that currently only humans can. Also, any of the quite different technologies that originate in that research field.

Consciousness

A term with many definitions. Functional or access consciousness refers roughly to there being global representations of information within a mind. Phenomenal consciousness refers to someone (or something) having experiences. There are many other ways of dividing up the concept of consciousness, and an active debate about which types of consciousness imply the others.

Diffusion model

A neural network that can generate images or sometimes other types of content such as 3D shapes. A diffusion model is trained to denoise existing noisy images in small steps. When this is done on complete noise, new images can be generated that correspond to prompts.

Ethology

The study of animal behavior. Modern ethology studies animal behavior and cognition in the context of the niche the animal's species is adapted to.

Evolutionary algorithm

An algorithm inspired by Darwinian evolution. In an evolutionary algorithm, a population of solutions is maintained, the best solutions (according to a fitness function) are kept, and the less good ones are replaced by mutated offspring of the best ones.

Fine-tuning

Further training of a model after pretraining. Typically, fine-tuning occurs on a smaller and more specific data set, making the model more suited to some specific task.

Foundation model

A typically very large machine learning model that can be used for multiple tasks, by either prompting or fine-tuning or something else. For example, a diffusion model such as Stable Diffusion can be used for image generation but also for editing images and changing their visual styles.

Generative adversarial network (GAN)

A type of neural network that can generate images. When training a GAN, two networks are used: the generator and the discriminator. The generator tries to generate new images that can pass as real according to the discriminator, and the discriminator tries to tell real images apart from the generated ones.

Gradient descent

The most common method for training neural networks. Compares the output of the neural network with the correct output (the label) and adjusts the weights of the neural network so that the output moves slightly toward the label. Doing this many times on a large number of data instances leads to the network gradually learning to produce the correct outputs.

Intelligence

A somewhat nebulous concept, variously understood as our ability to do things, our knowledge and understanding of the world, and our capacity to learn new things.

Intelligence quotient (IQ)

A score measured by one of several tests called intelligence tests. Depending on which particular test is used, the score might emphasize puzzle-solving ability, vocabulary, linguistic and numerical skills, or factual knowledge. IQ scores are usually normalized and age adjusted so that the mean score is 100 for any age group.

Large language model (LLM)

A large neural network, typically a transformer, that has been trained on an enormous amount of text (and often also program code). The training simply tasks the LLM with predicting the next token (letter or part of word) in a text. The trained LLM can then be used to generate text, and it turns out that lots of tasks can be phrased as text generation problems.

Model

An overloaded term. In the context of machine learning, a model is something that is trained from data (or interactions) and then used to produce predictions or actions. Examples include a neural network, decision tree, or support vector machine.

Neural network

A type of machine learning model vaguely inspired by biological brains. In practice, a neural network takes the form of several large matrices of numbers that are multiplied in various ways. Both the input and the output of the network take the form of numbers that are inputted into text, pixels, or something else. The neural network can be trained by gradient descent or some other algorithm, such as an evolutionary algorithm.

Niche

In ecology, the role an organism plays in a community. This includes the physical environment it has and how it interacts with other species. Through evolution, a species adapts to its particular niche (or niches). This includes the cognitive and behavioral aspects of the species as well as the physical aspects.

Pretraining

The initial training of a model, which creates concepts and a knowledge base that can later be exploited in fine-tuning.

Prompt

The text given as an input to an LLM or an image generator such as a diffusion model.

Reinforcement learning

A type of machine learning where an agent learns by interacting with the world and receiving (positive or negative) rewards for what it does; it is basically a formalization of learning by trial and error.

Self-supervised learning

A type of machine learning where a model predicts data without labels. For example, the model might predict the next character in a text, or learn to reconstruct noisy images. Self-supervised learning can be seen as supervised learning if one assumes that the input data are their own labels.

Supervised learning

A type of machine learning where a model is trained to predict data labels based on inputs (or features). It requires a data set with labeled instances. For example, an instance might have inputs corresponding to demographic information about a person, and the label might be whether that person would buy a used car or not; supervised learning can then be used to learn to predict who will buy a used car.

Transformer

A neural network architecture that underlies most LLM architectures. Transformers are sequential models that take sequences of “tokens” (which can be parts of words, parts of images, or something else) as inputs and produce new tokens. They use a technique called self-attention to correlate tokens in very different parts of the input.

Tree search

A class of search methods that start from an original state and search the space of possible future states resulting from taking actions in a systematic way. Classic AI planning builds on tree search algorithms, as do classic game-playing algorithms.

Chapter 1

1. Bill Gates, "The Age of AI Has Begun," *GatesNotes* (blog), March 21, 2023, <https://www.gatesnotes.com/The-Age-of-AI-Has-Begun>.
2. Glenn Hunter, "Exclusive Q&A: John Carmack's 'Different Path' to Artificial General Intelligence," *Dallas Innovates*, February 2, 2023, <https://dallasinnovates.com/exclusive-qa-john-carmacks-different-path-to-artificial-general-intelligence/>.
3. Murray Shanahan, *The Technological Singularity* (Cambridge, MA: MIT Press, 2015).
4. Ben Goertzel, "Artificial General Intelligence: Concept, State of the Art, and Future Prospects," *Journal of Artificial General Intelligence* 5, no. 1 (2014): 1.
5. Shane Legg and Marcus Hutter, "A Collection of Definitions of Intelligence," *Frontiers in Artificial Intelligence and Applications* 157 (2007): 17.

Chapter 2

1. M. Mitchell Waldrop, *The Dream Machine* (Toronto: Stripe Press, 2018).
2. Alfred North Whitehead and Bertrand Russell, *Principia Mathematica* (Cambridge: Cambridge University Press, 1910).
3. Allen Newell and Herbert Simon, "The Logic Theory Machine—a Complex Information Processing System," *IRE Transactions on Information Theory* 2, no. 3 (1956): 61–79.
4. T. Anthony Marsland, "A Review of Game-Tree Pruning," *ICGA Journal* 9, no. 1 (1986): 3–19.
5. Nathan Ensmenger, "Is Chess the Drosophila of Artificial Intelligence? A Social History of an Algorithm," *Social Studies of Science* 42, no. 1 (2012): 5–30.
6. Murray Campbell, A. Joseph Hoane Jr., and Feng-hsiung Hsu, "Deep Blue," *Artificial Intelligence* 134, nos. 1–2 (2002): 57–83.
7. Garry Kasparov, *Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins* (London: John Murray, 2017).
8. David Silver et al., "Mastering the Game of Go with Deep Neural Networks and Tree Search," *Nature* 529, no. 7587 (2016): 484–489.
9. Elizabeth Gibney, "Go Players React to Computer Defeat," *Nature*, January 2016, 24.
10. The Stockfish chess engine was created by Tord Romstad, Marco Costalba, and Joona Kiiski. It is freely available at stockfishchess.org.

11. David A. Forsyth and Jean Ponce, *Computer Vision: A Modern Approach* (Englewood Cliffs, NJ: Prentice Hall, 2002).
12. Jia Deng et al., "ImageNet: A Large-Scale Hierarchical Image Database," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (IEEE, 2009), 248–255.
13. Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," *Communications of the ACM* 60, no. 6 (2017): 84–90.
14. Chris M. Bishop, "Neural Networks and Their Applications," *Review of Scientific Instruments* 65, no. 6 (1994): 1803–1832.
15. Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, "Deep Learning," *Nature* 521, no. 7553 (2015): 436–444.
16. Simon Kornblith, Jonathon Shlens, and Quoc V. Le, "Do Better ImageNet Models Transfer Better?," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (IEEE, 2019), 2661–2671.
17. John Nickolls and William J. Dally, "The GPU Computing Era," *IEEE Micro* 30, no. 2 (2010): 56–69.

Chapter 3

1. Edwin G. Boring, "Intelligence as the Tests Test It," *New Republic* 36 (1923): 35–37.
2. Robert M. Thorndike, "The Early History of Intelligence Testing," in *Contemporary Intellectual Assessment: Theories, Tests, and Issues*, ed. D. P. Flanagan, J. L. Genshaft, and P. L. Harrison (New York: Guilford Press, 1997), 3–16.
3. Gale H. Roid and Mark Pomplun, *The Stanford-Binet Intelligence Scales* (New York: Guilford Press, 2012).
4. Kevin S. McGrew, "The Cattell-Horn-Carroll Theory of Cognitive Abilities: Past, Present, and Future," in *Contemporary Intellectual Assessment: Theories, Tests, and Issues*, 2nd ed., ed. D. P. Flanagan and P. L. Harrison (New York: Guilford Press, 2005), 136–181.
5. W. Joel Schneider and Kevin S. McGrew, "The Cattell-Horn-Carroll Theory of Cognitive Abilities," in *Contemporary Intellectual Assessment: Theories, Tests, and Issues*, 4th ed., ed. Dawn P. Flanagan and Erin M. McDonough (New York: Guilford Press, 2018), 73–163.
6. John L. Horn and Raymond B. Cattell, "Age Differences in Fluid and Crystallized Intelligence," *Acta Psychologica* 26 (1967): 107–129.
7. Kevin J. Mitchell, *Innate: How the Wiring of Our Brains Shapes Who We Are* (Princeton, NJ: Princeton University Press, 2020).

8. Nassim Nicholas Taleb, "Statistical Consequences of Fat Tails: Real World Preasymptotics, Epistemology, and Applications," arXiv preprint, 2020, arXiv: 2001.10488.
9. Antonio R. Damasio, *Descartes' Error* (New York: Random House, 2006).
10. Howard E. Gardner, *Frames of Mind: The Theory of Multiple Intelligences* (New York: Basic Books, 2011).
11. George Miller, "Varieties of Intelligence," *New York Times Book Review*, December 25, 1983, 5–25.
12. George Graham, "Behaviorism," *Stanford Encyclopedia of Philosophy* (2000), <https://plato.stanford.edu/entries/behaviorism/>.
13. Ivan Petrovich Pavlov, *Conditioned Reflexes: An Investigation of the Physiological Activity of the Cerebral Cortex*, ed. and trans. G. V. Anrep (1927; repr., New York: Boyer, 1960).
14. Burrhus Frederic Skinner, "'Superstition' in the Pigeon," *Journal of Experimental Psychology* 38, no. 2 (1948): 168.
15. John Garcia and Robert A. Koelling, "Relation of Cue to Consequence in Avoidance Learning," *Psychonomic Science* 4 (1966): 123–124.
16. Frans de Waal, *Are We Smart Enough to Know How Smart Animals Are?* (New York: W. W. Norton, 2016).
17. Christopher Flynn Martin et al., "Chimpanzee Choice Rates in Competitive Games Match Equilibrium Game Theory Predictions," *Scientific Reports* 4, no. 1 (2014): 1–6.
18. Walter T. Herbranson, "Pigeons, Humans, and the Monty Hall Dilemma," *Current Directions in Psychological Science* 21, no. 5 (2012): 297–301.
19. Jan K. Tornick et al., "Clark's Nutcrackers (*Nucifraga columbiana*) Use Gestures to Identify the Location of Hidden Food," *Animal Cognition* 14 (2011): 117–125.
20. Ed Yong, *An Immense World: How Animal Senses Reveal the Hidden Realms around Us* (Knopf Canada, 2022).
21. Peter Godfrey-Smith, *Other Minds: The Octopus, the Sea, and the Deep Origins of Consciousness* (New York: Farrar, Straus and Giroux, 2016).
22. John Alcock, *Animal Behavior: An Evolutionary Approach* (Sunderland, MA: Sinauer Associates, 2009).
23. Juliane Bräuer et al., "Old and New Approaches to Animal Cognition: There Is Not 'One Cognition,'" *Journal of Intelligence* 8, no. 3 (2020): 28.

Chapter 4

1. Alan Turing, "Computing Machinery and Intelligence," *Mind* 59, no. 236 (1950): 433.

2. Joseph Weizenbaum, "ELIZA—a Computer Program for the Study of Natural Language Communication between Man and Machine," *Communications of the ACM* 9, no. 1 (1966): 36–45.
3. Richard S. Sutton and Andrew G. Barto, *Reinforcement Learning: An Introduction* (Cambridge, MA: MIT Press, 1998).
4. D. H. Wolpert and W. G. Macready, "Coevolutionary Free Lunches," *IEEE Transactions on Evolutionary Computation* 9, no. 6 (2005): 721–735.
5. Stavros P. Adam et al., "No Free Lunch Theorem: A Review," in *Approximation and Optimization: Algorithms, Complexity and Applications*, ed. Ioannis C. Demetriou and Panos M. Pardalos (Cham, Switzerland: Springer, 2019), 57–82.
6. Shane Legg and Marcus Hutter, "Universal Intelligence: A Definition of Machine Intelligence," *Minds and Machines* 17 (2007): 391–444.
7. Marcus Hutter, *Universal Artificial Intelligence: Sequential Decisions Based on Algorithmic Probability* (Berlin: Springer Science & Business Media, 2004).
8. Yinon M. Bar-On, Rob Phillips, and Ron Milo, "The Biomass Distribution on Earth," *Proceedings of the National Academy of Sciences* 115, no. 25 (2018): 6506–6511.
9. François Chollet, "On the Measure of Intelligence," arXiv preprint, 2019, arXiv:1911.01547.
10. Bas Steunebrink, Jerry Swan, and Eric Nivel, "The Holon System: Artificial General Intelligence as 'Work on Command,'" *Proceedings of the Third International Workshop on Self-Supervised Learning* 192 (2022): 120–126.

Chapter 5

1. Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (New York: Oxford University Press, 2014).
2. Gary Marcus and Ernest Davis, *Rebooting AI: Building Artificial Intelligence We Can Trust* (New York: Vintage, 2019).

Chapter 6

1. Stuart J. Russell, *Artificial Intelligence: A Modern Approach* (London: Pearson Education, 2010).
2. James Moor, "The Dartmouth College Artificial Intelligence Conference: The Next Fifty Years," *AI Magazine* 27, no. 4 (2006): 87–91.
3. Bruce G. Buchanan and Reid G. Smith, "Fundamentals of Expert Systems," *Annual Review of Computer Science* 3, no. 1 (1988): 23–58.
4. Douglas B. Lenat et al., "Cyc: Toward Programs with Common Sense," *Communications of the ACM* 33, no. 8 (1990): 30–49.

5. Nils J. Nilsson, ed., *Shakey the Robot*, Technical Note 323 (Menlo Park, CA: AI Center, SRI International, 1984).
6. Ronald C. Arkin, *Behavior-Based Robotics* (Cambridge, MA: MIT Press, 1998).
7. Stevan Harnad, "The Symbol Grounding Problem," *Physica D: Nonlinear Phenomena* 42, nos. 1–3 (1990): 335–346.
8. Thomas M. Mitchell, *Machine Learning* (New York: McGraw Hill, 1997).
9. David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams, "Learning Representations by Back-Propagating Errors," *Nature* 323, no. 6088 (1986): 533–536.
10. Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep Learning* (Cambridge, MA: MIT Press, 2016).
11. Wenyi Zhao et al., "Face Recognition: A Literature Survey," *ACM Computing Surveys (CSUR)* 35, no. 4 (2003): 399–458.
12. Jonathan Masci et al., "Stacked Convolutional Auto-encoders for Hierarchical Feature Extraction," in *Artificial Neural Networks and Machine Learning—ICANN 2011: 21st International Conference on Artificial Neural Networks, Espoo, Finland, June 14–17, 2011, Proceedings*, part I (Berlin: Springer, 2011), 52–59.
13. Felipe Codevilla et al., "End-to-End Driving via Conditional Imitation Learning," in *2018 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2018), 4693–4700.
14. Dean A. Pomerleau, "Alvin: An Autonomous Land Vehicle in a Neural Network," *Advances in Neural Information Processing Systems* 1 (1988): 305–313.
15. Arthur L. Samuel, "Some Studies in Machine Learning Using the Game of Checkers," *IBM Journal of Research and Development* 3, no. 3 (1959): 210–229.
16. Richard S. Sutton and Andrew G. Barto, *Reinforcement Learning: An Introduction* (Cambridge, MA: MIT Press, 2018).
17. Gerald Tesauro, "TD-Gammon, a Self-Teaching Backgammon Program, Achieves Master-Level Play," *Neural Computation* 6, no. 2 (1994): 215–219.
18. Volodymyr Mnih et al., "Human-Level Control through Deep Reinforcement Learning," *Nature* 518, no. 7540 (2015): 529–533.
19. Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter, "Neural Architecture Search: A Survey," *Journal of Machine Learning Research* 20, no. 1 (2019): 1997–2017.
20. Niels Justesen et al., "Deep Learning for Video Game Playing," *IEEE Transactions on Games* 12, no. 1 (2019): 1–20.
21. Chiyuan Zhang et al., "A Study on Overfitting in Deep Reinforcement Learning," arXiv preprint, 2018, arXiv:1804.06893.

22. Gregory Hornby et al., “Automated Antenna Design with Evolutionary Algorithms” (paper presented at *Space 2006*, San Jose, CA, September 19–21, 2006); Julian Togelius et al., “Search-Based Procedural Content Generation: A Taxonomy and Survey,” *IEEE Transactions on Computational Intelligence and AI in Games* 3, no. 3 (2011): 172–186; Razieh Saremi et al., “An Evolutionary Algorithm for Task Scheduling in Crowdsourced Software Development,” arXiv preprint, 2021, arXiv:2107.02202; Stefano Nolfi and Dario Floreano, *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines* (Cambridge, MA: MIT Press, 2000).
23. Marvin L. Minsky, “Logical versus Analogical or Symbolic versus Connectionist or Neat versus Scruffy,” *AI Magazine* 12, no. 2 (1991): 34–51.
24. David Silver et al., “Mastering the Game of Go with Deep Neural Networks and Tree Search,” *Nature* 529, no. 7587 (2016): 484–489.

Chapter 7

1. Mark A. Kramer, “Nonlinear Principal Component Analysis Using Autoassociative Neural Networks,” *Aiche Journal* 37, no. 2 (1991): 233–243.
2. Ian Goodfellow et al., “Generative Adversarial Networks,” *Communications of the ACM* 63, no. 11 (2020): 139–144.
3. Tero Karras et al., “Analyzing and Improving the Image Quality of Stylegan,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (IEEE, 2020), 8110–8119.
4. Aditya Ramesh et al., “Zero-Shot Text-to-Image Generation,” in *International Conference on Machine Learning* (PMLR, 2021), 8821–8831; Jascha Sohl-Dickstein et al., “Deep Unsupervised Learning Using Nonequilibrium Thermodynamics,” in *International Conference on Machine Learning* (PMLR, 2015), 2256–2265; Jonathan Ho, Ajay Jain, and Pieter Abbeel, “Denoising Diffusion Probabilistic Models,” *Advances in Neural Information Processing Systems* 33 (2020): 6840–6851.
5. Robin Rombach et al., “High-Resolution Image Synthesis with Latent Diffusion Models,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (IEEE, 2022), 10684–10695.
6. Sébastien Bubeck et al., “Sparks of Artificial General Intelligence: Early Experiments with GPT-4,” arXiv preprint, 2023, arXiv:2303.12712.
7. C. E. Shannon, “A Mathematical Theory of Communication,” *Bell System Technical Journal* 27, no. 3 (1948): 379–423.
8. Geoffrey E. Hinton, *Connectionist Symbol Processing* (Cambridge, MA: MIT Press, 1991).

9. Jeffrey L. Elman, "Finding Structure in Time," *Cognitive Science* 14, no. 2 (1990): 179–211.
10. Sepp Hochreiter and Jürgen Schmidhuber, "Long Short-Term Memory," *Neural Computation* 9, no. 8 (1997): 1735–1780.
11. Ashish Vaswani et al., "Attention Is All You Need," *Advances in Neural Information Processing Systems* 30 (2017).
12. Alec Radford et al., "Language Models Are Unsupervised Multitask Learners," *OpenAI Blog* 1, no. 8 (2019): 9.
13. Tom Brown et al., "Language Models Are Few-Shot Learners," *Advances in Neural Information Processing Systems* 33 (2020): 1877–1901.
14. Jason Wei et al., "Finetuned Language Models Are Zero-Shot Learners," arXiv preprint, 2021, arXiv:2109.01652.
15. Long Ouyang et al., "Training Language Models to Follow Instructions with Human Feedback," *Advances in Neural Information Processing Systems* 35 (2022): 27730–27744.
16. OpenAI, "GPT-4 Technical Report," arXiv preprint, 2023, arXiv:2303.08774.
17. Michael Ahn et al., "Do as I Can, Not as I Say: Grounding Language in Robotic Affordances," arXiv preprint, 2022, arXiv:2204.01691.
18. Toran Bruce Richards, "Auto-GPT: An Autonomous GPT-4 Experiment" (2023), <https://github.com/Significant-Gravitas/AutoGPT>.
19. Sébastien Bubeck et al., "Sparks of Artificial General Intelligence: Early Experiments with GPT-4," arXiv preprint, 2023, arXiv:2303.12712.
20. R. Thomas McCoy et al., "Embers of Autoregression: Understanding Large Language Models through the Problem They Are Trained to Solve," arXiv preprint, 2023, arXiv:2309.13638.

Chapter 8

1. Richard Dawkins, *The Blind Watchmaker: Why the Evidence of Evolution Reveals a Universe without Design* (New York: W. W. Norton, 1996).
2. Agoston E. Eiben and James E. Smith, *Introduction to Evolutionary Computing* (Berlin: Springer, 2015).
3. Christopher G. Langton, ed., *Artificial Life: An Overview* (Cambridge, MA: MIT Press, 1997).
4. Mark A. Bedau, "Artificial Life," in *Philosophy of Biology* (Amsterdam: North-Holland, 2007), 585–603.
5. Nils Aall Barricelli, "Numerical Testing of Evolution Theories: Part I; Theoretical Introduction and Basic Tests," *Acta Biotheoretica* 16, nos. 1–2 (1962): 69–98.

6. Thomas S. Ray, "Evolution and Optimization of Digital Organisms," in *Scientific Excellence in Supercomputing: The IBM 1990 Contest Prize Papers*, ed. K. R. Billingsley, E. Derohanes, and H. Brown III (Athens, GA: Baldwin Press, University of Georgia, 1991).
7. Richard Dawkins and John Richard Krebs, "Arms Races between and within Species," *Proceedings of the Royal Society of London B: Biological Sciences* 205, no. 1161 (1979): 489–511.
8. W. Daniel Hillis, "Co-evolving Parasites Improve Simulated Evolution as an Optimization Procedure," *Physica D: Nonlinear Phenomena* 42, nos. 1–3 (1990): 228–234.
9. Dario Floreano and Stefano Nolfi, "God Save the Red Queen! Competition in Co-evolutionary Robotics," in *Proceedings of the 2nd Annual Conference on Genetic Programming* (Burlington, MA: Morgan Kaufmann, 1997).
10. Kenneth O. Stanley and Risto Miikkulainen, "Competitive Coevolution through Evolutionary Complexification," *Journal of Artificial Intelligence Research* 21 (2004): 63–100.
11. Christopher D. Rosin and Richard K. Belew, "New Methods for Competitive Coevolution," *Evolutionary Computation* 5, no. 1 (1997): 1–29.
12. Julian Togelius, "AI Researchers, Video Games Are Your Friends!," in *Computational Intelligence: International Joint Conference, IJCCI 2015 Lisbon, Portugal, November 12–14, 2015, Revised Selected Papers* (Berlin: Springer International Publishing, 2017), 3–18.
13. Julian Togelius, *Playing Smart: On Games, Intelligence, and Artificial Intelligence* (Cambridge, MA: MIT Press, 2019).
14. Ralph Koster, *Theory of Fun for Game Design* (Scottsdale, AZ: Paraglyph Press, 2005).
15. Frank Lantz et al., "Depth in Strategic Games," in *AAAI Workshops* (Washington, DC: AAAI, 2017).
16. Volodymyr Mnih et al., "Human-Level Control through Deep Reinforcement Learning," *Nature* 518, no. 7540 (2015): 529–533.
17. Chiyuan Zhang et al., "A Study on Overfitting in Deep Reinforcement Learning," arXiv preprint, 2018, arXiv:1804.06893.
18. Karl Cobbe et al., "Quantifying Generalization in Reinforcement Learning," in *International Conference on Machine Learning* (PMLR, 2019), 1282–1289.
19. Diego Perez-Liebana et al., "General Video Game AI: A Multitrack Framework for Evaluating Agents, Games, and Content Generation Algorithms," *IEEE Transactions on Games* 11, no. 3 (2019): 195–214.
20. Niels Justesen et al., "Illuminating Generalization in Deep Reinforcement Learning through Procedural Level Generation," *AAAI Workshop on Reinforcement Learning in Games* (2018).

21. Noor Shaker, Julian Togelius, and Mark J. Nelson, *Procedural Content Generation in Games* (Berlin: Springer, 2016).
22. Cameron Browne and Frederic Maire, "Evolutionary Game Design," *IEEE Transactions on Computational Intelligence and AI in Games* 2, no. 1 (2010): 1–16.
23. Julian Togelius and Jurgen Schmidhuber, "An Experiment in Automatic Game Design," in *2008 IEEE Symposium on Computational Intelligence and Games* (IEEE, 2008), 111–118; Thorbjørn S. Nielsen et al., "Towards Generating Arcade Game Rules with VGDL," in *2015 IEEE Conference on Computational Intelligence and Games (CIG)* (IEEE, 2015), 185–192.
24. Jakob Bauer et al., "Human-Timescale Adaptation in an Open-Ended Task Space," arXiv preprint, 2023, arXiv:2301.07608.

Chapter 9

1. Anita Avramides, "Other Minds," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta (Winter 2020 edition), <https://plato.stanford.edu/entries/other-minds/>.
2. David J. Chalmers, *The Conscious Mind: In Search of a Fundamental Theory* (New York: Oxford University Press, 1997).
3. Thomas Nagel, "What Is It Like to Be a Bat?," *Philosophical Review* 83, no. 4 (1974): 435–450.
4. Michael A. Cohen and Daniel C. Dennett, "Consciousness Cannot Be Separated from Function," *Trends in Cognitive Sciences* 15, no. 8 (2011): 358–364.
5. John Jamieson Carswell Smart, "The Content of Physicalism," *Philosophical Quarterly* 28, no. 113 (1978): 339–341.
6. Ned J. Block, "Functionalism," in *Studies in Logic and the Foundations of Mathematics*, vol. 104 (Amsterdam: Elsevier, 1982), 519–539.
7. David J. Chalmers, "Absent Qualia, Fading Qualia, Dancing Qualia," in *Conscious Experience*, ed. Thomas Metzinger (Paderborn: Ferdinand Schöningh, 1995), 309–328.
8. David J. Chalmers, "Could a Large Language Model Be Conscious?," *Boston Review*, August 9, 2023.
9. Patrick Butlin et al., "Consciousness in Artificial Intelligence: Insights from the Science of Consciousness," arXiv preprint, 2023, arXiv:2308.08708.

Chapter 10

1. Irving John Good, "Speculations Concerning the First Ultraintelligent Machine," in *Advances in Computers*, vol. 6 (Amsterdam: Elsevier, 1966), 31–88.
2. Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (New York: Oxford University Press, 2014).

3. David J. Chalmers, "The Singularity: A Philosophical Analysis," *Journal of Consciousness Studies* 17, nos. 9–10 (2010): 7–65.
4. Stephen Hawking, *Brief Answers to the Big Questions* (New York: Bantam, 2018).
5. "Statement on AI Risk," Center for AI Safety, accessed March 21, 2024, <https://www.safe.ai/statement-on-ai-risk>.
6. Nick Bostrom, "Ethical Issues in Advanced Artificial Intelligence," in *Cognitive, Emotive and Ethical Aspects of Decision Making in Humans and in Artificial Intelligence*, vol. 2, ed. I. Smit et al. (Tecumseh, Ontario: International Institute of Advanced Studies in Systems Research and Cybernetics, 2003), 12–17.

Chapter 11

1. Peter Stone et al., "Artificial Intelligence and Life in 2030: The One Hundred Year Study on Artificial Intelligence," arXiv preprint, 2022, arXiv:2211.06318.
2. Iain M. Banks, *The Player of Games* (New York: Macmillan, 1988).
3. Hans Moravec, *Mind Children: The Future of Robot and Human Intelligence* (Cambridge, MA: Harvard University Press, 1988).
4. James Bessen, *Learning by Doing: The Real Connection between Innovation, Wages, and Wealth* (New Haven, CT: Yale University Press, 2015).
5. Jacob Goldstein, "How the Electronic Spreadsheet Revolutionized Business," NPR, February 27, 2015, <https://www.npr.org/2015/02/27/389585340/how-the-electronic-spreadsheet-revolutionized-business>.
6. Jennifer Smith, "Self-Driving Technology Threatens Nearly 300,000 Trucking Jobs, Report Says," *Wall Street Journal*, September 4, 2018.
7. Rick Rubin, "Star Power/Hide and Seek/The Guru," interview by Anderson Cooper, *60 Minutes*, CBS, January 15, 2023.
8. Tomas Mikolov et al., "Efficient Estimation of Word Representations in Vector Space," arXiv preprint, 2013, arXiv:1301.3781.
9. Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan, "Semantics Derived Automatically from Language Corpora Contain Human-Like Biases," *Science* 356, no. 6334 (2017): 183–186.
10. Luhang Sun et al., "Smiling Women Pitching Down: Auditing Representational and Presentational Gender Biases in Image Generative AI," arXiv preprint, 2023, arXiv:2305.10566.
11. Kate Crawford, "The Trouble with Bias" (keynote address at Conference on Neural Information Processing Systems, December 5, 2017, https://youtu.be/fMym_BKWQzk?si=RdPKdDA8xuIxnKY1).
12. Jiawen Deng et al., "Recent Advances towards Safe, Responsible, and Moral Dialogue Systems: A Survey," arXiv preprint, 2023, arXiv:2302.09270.

13. Tony Sun et al., "Mitigating Gender Bias in Natural Language Processing: Literature Review," arXiv preprint, 2019, arXiv:1906.08976.
14. Gary Marcus, "Hoping for the Best as AI Evolves," *Communications of the ACM* 66, no. 4 (2023): 6–7.
15. Kailey Huang, "Why Pope Francis Is the Star of A.I.-Generated Photos," *New York Times*, April 8, 2023, <https://www.nytimes.com/2023/04/08/technology/ai-photos-pope-francis.html>.
16. Philip Bontrager et al., "Deepmasterprints: Generating Masterprints for Dictionary Attacks via Latent Variable Evolution," in *2018 IEEE 9th International Conference on Biometrics Theory, Applications and Systems (BTAS)* (IEEE, 2018).

Chapter 12

1. Alan Turing, "Computing Machinery and Intelligence," *Mind* 59, no. 236 (1950): 433.

FURTHER READING

Chalmers, David J. *Reality+: Virtual Worlds and the Problems of Philosophy*. London: Penguin, 2022.

Chollet, Francois. *Deep Learning with Python*. 2nd ed. Shelter Island, NY: Manning, 2021.

Chollet, François. "On the Measure of Intelligence." arXiv preprint. 2019. <https://arxiv.org/abs/1911.01547>.

De Waal, Frans. *Are We Smart Enough to Know How Smart Animals Are?* New York: W. W. Norton, 2016.

Li, Fei-Fei. *The Worlds I See: Curiosity, Exploration, and Discovery at the Dawn of AI*. New York: Flatiron Books, 2023.

Marcus, Gary, and Ernest Davis. *Rebooting AI: Building Artificial Intelligence We Can Trust*. New York: Pantheon, 2019.

Mitchell, Melanie. *Artificial Intelligence: A Guide for Thinking Humans*. London: Penguin, 2019.

Russell, Stuart J., and Peter Norvig. *Artificial Intelligence: A Modern Approach*. 4th ed. London: Pearson, 2020.

Wooldridge, Michael. *A Brief History of Artificial Intelligence: What It Is, Where We Are, and Where We Are Going*. New York: Flatiron Books, 2021.

INDEX

- Abstraction and Reasoning Corpus (ARC), 66, 68
Adaptive Agent (program), 139
Aerosmith, 177
Age bias, 179
AIXI, 62, 72
AlexNet, 26–27, 29, 31, 89, 97
Aliens (film), 75
AlphaGo, 2, 23, 31, 97, 99
Animal intelligence, 5, 31, 41–51, 53, 65, 146, 188
Ants, 9, 46, 157
Artificial life, 129
Atari, 93, 135–136
Attention, neural network architecture. *See* Transformer, neural network architecture
Autocomplete, 1, 190
Autoencoders, 102–104
AutoGPT (program), 121
Avengers: Age of Ultron (film), 156

Babbage, Charles, 15
Backgammon, 92
Bacteria, 65
Banks, Iain M., 76, 175. *See also* Culture series
Barricelli, Nils Aall, 129
Bats, 45
Beastie Boys, 177
Behaviorism, 41–43, 57
Bias, social, 122, 178–180
Binet, Alfred, 33–34
Biometrics, 89, 182

Board game playing, 20–24, 28, 84, 92, 97, 99, 138
Boring, Edwin, 33
Bostrom, Nick, 74, 157, 163
Boulder Dash (video game), 136
Browne, Cameron, 138

Carlsen, Magnus, 123
Carmack, John, 3
Cash, Johnny, 177
Cats, 41, 45, 50, 58
Cattell-Horn-Carroll (CHC) theory, 35–38, 40, 66
Chalmers, David, 146–151, 157
ChatGPT, 55, 118–120, 185. *See also* GPT
Checkers, 23, 91–92
Chess, 2, 5–6, 20–28, 55, 123, 186, 200
Chollet, François, 65–66
Churchill, Winston, 125
Clarke, Arthur C., 156
Clippy, 79–80
Coevolution, 131–133
Collective intelligence, 43, 74, 78, 82
Competitive coevolution, 132–133
Consciousness, 9, 74–76, 118, 145–155
Cooper, Anderson, 177
Cooperative coevolution, 131–132
Corvids, 44, 48–49
Creativity, 22, 38
Crystallized intelligence, 5, 35–36, 66

- Culture series (Banks), 76–77, 168–169, 175
- Cybersecurity, 165, 178, 182–183, 196
- Cyc project, 85

- Dall-E (program), 104, 190
- Damasio, Antonio, 38
- Decision trees, 87
- Deep Blue (program), 21–23, 28, 31, 83–84, 123
- Deep learning. *See* Neural networks
- DeepMind, 23, 97, 139
- Dennett, Daniel, 147
- De Waal, Frans, 44
- Diffusion models, 104, 114
- Dogs, 41, 42, 45, 49, 175, 188
- Donkey Kong* (video game), 136
- Doom* (video game), 3

- Eliza (program), 54–55
- Embodiment, 71, 76, 79–80, 168, 189
- Employment, 10, 58, 172–178
- Environment, for agents, 57–66, 70, 95–96, 128, 133–134, 138, 140–142, 188, 192
- Ethology, 43
- Evolution. *See also* Coevolution
 - algorithmic, 95, 99, 127–134, 139–140, 158
 - natural, 43–49, 63–65, 94, 127–134, 191
- Ex Machina* (film), 75

- Features, machine learning, 90, 102, 138
- Fine-tuning, 27, 114, 116, 119, 122

- Fitness, evolutionary, 64–65, 94–96, 127, 131–133
- Fluid intelligence, 5, 35–36, 66
- Foundation* (TV show), 72
- Foundation models, 8, 105, 114, 116, 121, 123–126, 179–180
- Francis (pope), 105, 181
- Frogger* (video game), 136
- Functionalism, 148–150

- Gardner, Howard, 40
- Gates, Bill, 3, 158
- Gender bias, 178–180
- Generative adversarial network (GAN), 103–104
- g* factor, 7, 35–37, 41
- Go (board game), 2, 22–24, 31, 97, 99, 200
- Goertzel, Ben, 3
- GOF AI (good old-fashioned AI). *See* Symbolic AI
- Golem, 15
- Good, Irving John, 156
- Google, 23, 79, 92, 118, 121–122, 139
- GPT
 - ChatGPT, 55, 118–120, 185
 - GPT-2, 113–116
 - GPT-3, 116–118
 - GPT-4, 8, 19, 120–126, 190
- GPU (graphics processing unit), 27, 89, 111, 160
- Gradient descent, 97
- Grand Theft Auto* (video game), 93
- Graphics card. *See* GPU

- Hades* (video game), 138
- Hawking, Stephen, 158
- Hephaestus, 15

- Hillis, Daniel, 132–133
 Honeybees, 77
 Hume, David, 38
 Hutter, Marcus, 61
- I, Robot* (film), 156
- ImageNet, 25, 27, 29, 89, 97
- Image recognition, 24–28, 31, 186
- Industrial Revolution, 170
- Instruction tuning, 118, 190
- Intelligence
 animal, 5, 31, 41–51, 53, 65, 146, 188
 collective, 43, 74, 78, 82
 fluid vs. crystallized, 5, 35–36, 66
 superintelligence, 9, 155–165
 universal, 7, 61–65, 94, 188
- Intelligence tests, 33–37
- IQ (intelligence quotient), 33–41, 66, 188
- Jay-Z, 177
- Job loss. *See* Employment
- Kasparov, Garry, 21
- Kubrick, Stanley, 156
- Labels, machine learning, 25, 27, 101–102
- Large language models (LLMs), 19, 56, 113–126, 140, 152–153, 160, 172, 180–182, 190–191, 195
 instruction tuning, 118, 190
 prompting, 8, 110–112, 114, 116–123, 125, 152–153, 190
 reinforcement learning from human feedback (RLHF), 119, 180
 tokens, 113, 122, 125
- Legg, Shane, 61
- Lem, Stanisław, 72, 78, 189
- Logic Theorist (program), 17–19, 22, 28–29, 83–84
- Long short-term memory (LSTM), 110–112
- Machine learning, 26–29, 64, 87–91, 115, 182, 186
 features, 90, 102, 138
 labels, 25, 27, 101–102
- Mantis shrimp, 77
- Mathematics, 17, 19, 24, 188
- Matrix, The* (film), 155
- Meta, 122
- Microsoft, 79, 125, 158, 172, 189
- Midjourney (program), 105, 181, 190
- Miéville, China, 78, 189
- Miller, George, 40
- Minimax, 20–23
- Misinformation, 180–182
- Mistral, 122
- Moving goalposts phenomenon, 28, 200
- Musk, Elon, 157
- Nagel, Thomas, 147
- Neats vs. scruffies, 96–98
- Neural networks, 8, 23, 26–31, 57, 66, 82, 87, 89–101, 109, 115–116, 123, 133–136, 158–159, 178, 186, 198
 autoencoders, 102–104
 diffusion models, 104, 114
 generative adversarial network (GAN), 103–104
 long short-term memory (LSTM), 110–112

- Neural networks (cont.)
 - recurrent, 110–111, 125
 - transformer architecture, 113
- Newell, Allen, 17, 19
- n-gram, 106–109, 112, 122
- “No free lunch” theorem, 7, 59, 61, 68, 70, 77
- No Man’s Sky* (video game), 138
- Octopuses, 45–50, 66, 188
- OpenAI, 8, 104, 113–116, 120–121, 157
- Open-ended learning, 8, 78, 99, 127–131, 139–141, 191–192
- Optimization, 29, 61, 159, 190
- Overfitting, 93, 136
- Pac-Man* (video game), 136
- Pavlov, Ivan Petrovich, 42
- Philosophical zombies, 147–148
- Physicalism, 148–150
- Pigeons, 42, 44, 57–58
- Pocket calculators, 16, 195
- Pre-training, 111, 114, 116
- Principia Mathematica* (Whitehead and Russell), 17–19
- Prompting, language modeling, 8, 110–112, 114, 116–123, 125, 152–153, 190
- Q*bert* (video game), 136
- Racial bias, 179
- Rage against the Machine, 177
- Recurrent neural networks (RNN), 110–111, 125
- Red Hot Chili Peppers, 177
- Reinforcement learning, 57–58, 87, 91–93, 119, 127, 135–140, 158, 180
- Reinforcement learning from human feedback (RLHF), 119, 180
- Rogue* (video game), 138
- Route planning, 1
- Rubin, Rick, 177
- Russell, Bertrand, 18
- Samuels, Arthur, 91
- SayCan (robot control system), 121
- Scruffies vs. neats, 96–98
- Search algorithm, 17, 21–23, 29, 31, 61, 84, 97, 99, 186, 190
- Sedol, Lee, 23
- Self-driving cars, 2, 90–91, 93–94, 170, 173–174
- Self-supervised learning, 8, 87, 99, 101–106, 112, 141, 190, 192
- Shakey (robot), 85–86
- Shanahan, Murray, 3
- Shannon, Claude, 106–107, 112
- Simon, Herbert, 17, 19
- Skinner, B. F., 42
- Snoop Dogg, 175
- Social bias, 122, 178–180
- Sokoban* (video game), 116
- Space Invaders* (video game), 136
- Spelunky* (video game), 138
- Stable Diffusion, 8, 104–105, 114, 179, 190
- StarCraft* (video game), 135
- Star Trek* (TV series), 1, 75, 79–80
- Star Wars franchise, 70, 76
- Street Fighter* (video game), 135
- Superintelligence, 9, 155–165

- Supervised learning, 61, 67, 69, 90–91, 97, 99, 101–102, 105
- Support vector machines, 87
- Symbol-grounding problem, 86
- Symbolic AI, 83–86, 97

- TD-Gammon (algorithm), 92
- Terminator, The* (film series), 1–2, 71, 74, 155, 162
- Tierra (program), 129
- Tokens, language modeling, 113, 122, 125
- Transformer, neural network architecture, 113
- Transformers* (film series), 156
- Tree (search), 20, 186
- Trotsky, Leon, 181
- Turing, Alan, 53–54, 201
- Turing test, 53–55, 65, 68, 201
- 2001: A Space Odyssey* (film), 2, 71, 74, 156

- Universal intelligence, 7, 61–65, 94, 188

- Video games, 1, 24, 57–58, 70, 72, 89, 92–94, 116, 134–139, 191, 200

- WALL-E* (film), 1–2, 71
- War Games* (film), 163
- Weizenbaum, Joseph, 54–55
- Westworld* (TV show), 72
- Whitehead, Alfred North, 18
- Word2vec (program), 178

- Yavalath* (board game), 138

JULIAN TOGELIUS is Associate Professor of Computer Science and Engineering at New York University, and cofounder of the game AI startup modl.ai. He works on artificial intelligence for video games, computational creativity, evolutionary computation, and applications of AI in fields from education to biometrics. He originally studied philosophy and psychology to try to understand the mind but then decided it was more fruitful to try to build minds to understand them.