

POCKET EINSTEIN

10 Short Lessons in
Artificial Intelligence & Robotics

First published in Great Britain in 2020
by Michael O'Mara Books Limited
9 Lion Yard
Tremadoc Road
London SW4 7NQ

Copyright © Michael O'Mara Books 2020

All rights reserved. You may not copy, store, distribute, transmit, reproduce or otherwise make available this publication (or any part of it) in any form, or by any means (electronic, digital, optical, mechanical, photocopying, recording or otherwise), without the prior written permission of the publisher. Any person who does any unauthorized act in relation to this publication may be liable to criminal prosecution and civil claims for damages.

A CIP catalogue record for this book is available from the British Library.

ISBN: 978-1-78929-216-9 in hardback print format
ISBN: 978-1-78929-217-6 in ebook format

www.mombooks.com

CONTENTS

Introduction

LESSON 1: A Journey of a Thousand Miles Begins with a Single Step

LESSON 2: Choose the Right Path

LESSON 3: We All Fall Down

LESSON 4: Find the Right Answer

LESSON 5: Understand Your World

LESSON 6: Change for the Better

LESSON 7: Communication is Key

LESSON 8: Re-imagine Reality

LESSON 9: Feel Better

LESSON 10: Know Yourself

Glossary

Further Reading

Index

INTRODUCTION

I grew up in the 1970s when there was no internet, no World Wide Web, and the very first affordable home computers were just emerging. In those days you had to be a serious geek like me to be into computers. Yes, I was that kid – shy at school, but prolific at home, building bizarre robots, programming early computers, writing simple computer games. I would lust after the latest computer as another child might desire a Lamborghini. So amazing, but so unaffordable! Artificial intelligence – getting my computers to think, to simulate biological behaviour and control my robots – was my childhood passion. But to the people around me, it was a passion that seemed as obscure as stamp collecting might today.

And then things changed. Quite dramatically. Today we live in a science-fiction story come true. Computers rule the world. Data floods from everything we do. Robots are building our products in factories. Our homes are computerized and we can talk to these digital home assistants, receiving detailed and coherent replies. Behind the scenes, artificial intelligence is making everything work. My obscure childhood passion is now not only mainstream, it's considered one of the most

important kinds of technology being created today.

You cannot live in the modern world without interacting with, or being impacted by, AI and robots. Every time you make a purchase, AIs are handling your money, checking for fraud, using your data to understand you better, and recommending new products to you. Every time you drive a car, AIs are helping the car to proceed safely, they're watching you from road cameras and automatically changing speed limits, they're detecting your licence plate and monitoring your movements. Every time you post something on social media, AIs may trawl through the text to understand your sentiment on specific topics. As you browse the internet and read news articles or blogs, AIs monitor your activity and try to please you by feeding you more of the content you prefer. Every time you take a photo, AIs adjust the camera settings and ensure the best possible picture is taken – and then can identify everyone in the picture for you. Face recognition, speech understanding, automatic bots that answer your questions online or by phone – all performed by AIs. Inside your home you have smart TVs, computerized fridges, washing machines, central heating, air-conditioning systems – all AI robotic devices. The world economy is managed by AIs, financial trading is performed by AIs, and decisions about whether you should or should not be accepted for financial products are made by AIs. Your future anti-viral and antibacterial drugs are being designed by AI. Your services of water, electricity, gas, mobile phone and internet connections are all adjusted by smart AI algorithms that try to optimize supply while minimizing waste. You interact with a thousand AIs a day and you are blissfully unaware of them all.

In this book I'll explain a little bit about how this has happened, how it works and what it means. This is a pocket guide, so I'm going to be brief. I won't bore you with detailed technical descriptions, I will not

explain every single AI technique, and I will not tell you about every AI pioneer. That would take a thousand books of this size, with more books needed every day (progress is fast!). Instead I'll take you on a short journey through this strange world of computers, robots and building brains. I'll try to point out some interesting sights along the way, and explain some of the fundamental ideas behind artificial intelligence and robotics. This journey may sometimes be a rollercoaster, for AI has its ups and downs. It has lived a surprisingly long life already, and suffered the pains of disappointment as well as the excitement of success. It is being created to change our world for the better, yet in some cases it is responsible for causing fundamental problems. Buckle up, and enjoy the ride!

Peter J. Bentley

01 A JOURNEY OF A THOUSAND MILES BEGINS WITH A SINGLE STEP

'I confidently predict that in the next ten or fifteen years something will emerge from the laboratory that is not dissimilar to the robot of science-fiction fame.'

CLAUDE SHANNON (1961)

Classical stone architecture and statues surround you. You walk through the cobbled streets of the pretty Greek island, admiring the view. The hot sun is now low in the sky, leaving a pleasant evening for a stroll around the town. The hustle and bustle of everyday life has faded away as the market stalls of fruits and fish are closed. There is just the sound of your own footsteps echoing from the ornate buildings. An unexpected movement catches your eye from the corner of the street. Yet there is nobody there. You look harder. The stone statue – it moved! You nervously walk over for a closer look. Its chest appears to rise and fall as

though it breathes. As you watch, its head turns left, then right. You realize that it's not the only one. All the statues on the streets around you seem to display some movement. They move their feet as if adjusting position, they move arms as though having some silent stone discussion. Are they slowly coming to life as night falls? Looking closely, you realize they all seem to have hidden mechanisms, cogs and wheels whirring. You're on an island of stone robots.


Ancient robots

This was the Greek island of Rhodes, as you may have found it 2,400 years ago, even before their giant statue, Colossus of Rhodes, was constructed. It was a remarkable island famed for its mechanical inventions, including life-sized automata made from marble. An Ancient Greek poet named Pindar visited Rhodes and wrote about his experience in a poem:

*The animated figures stand
Adorning every public street
And seem to breathe in stone, or
move their marble feet.*

It may seem inconceivable that before the Roman Empire in 400 BCE there was such ancient robotic technology. But many ancient examples are well documented. Powered by water or weights, there were mechanical lions that roared, metal birds that sang, and even mechanical people who played together in a band. King Solomon, who reigned from 970 to 931 BCE, was said to have had a golden lion that raised a foot to help him to his throne, and a mechanical eagle that placed his crown upon his head. Ancient Chinese texts tell the story of a mechanical man

presented to King Mu of Zhou (1023–957 BCE) by the ‘artificer’ Yan Shi. Archytas, founder of mathematical mechanics, philosopher and pal of Plato, who lived from 428 to 347 BCE, made a mechanical dove – a flying, steam-powered wooden robot bird. Hero of Alexandria (10–70 CE) wrote an entire book about his automaton inventions, and how hydraulics, pneumatics and mechanics could be used. Hero even created a programmable puppet show that used carefully measured lengths of thread that were pulled by a weight to trigger different events in his choreographed mechanical play.



This fascination with building mechanical life continued unabated through the Middle Ages. Countless inventors created mechanical marvels designed to entertain. By the eighteenth century this was taken to a new level by the inventors of automated factory machines, enabling the Industrial Revolution. Laborious jobs such as weaving that had always required skilled human craftspeople were suddenly replaced by astonishing steam-powered machines that could create finer cloth faster than ever before. As one set of jobs were lost, whole new industries were created as our massive machines needed constant care and maintenance.

The decades rolled past and our expertise in building machines

increased. Trains, automobiles, aeroplanes and sophisticated factories became commonplace. With an increasing reliance on automatic machines, the allure of robots and their similarity to living creatures only intensified, entering literature and movies. It is perhaps no coincidence that two of the very earliest science-fiction movies, *Metropolis* (1927) and *Frankenstein* (1931), tell the story of crazed inventors creating life.

By the twentieth century scientists were trying to understand life itself through making analogies. Perhaps, they thought, if we could make something that moved and thought like a living creature then we would learn the secrets behind life – understanding by making. This was the start of artificial intelligence (AI) and robotics as we know them today.

The birth of AI and robotics


One of the very earliest examples of an autonomous robot designed to help us understand living systems was built in the late 1940s by neurologist Grey Walter in Bristol, UK. Since they looked a little like electric tortoises, he named them Elmer and Elsie (ELECTRO MECHANICAL Robots, Light Sensitive). Grey Walter's robots were unique because they followed no specific program.

WILLIAM GREY WALTER (1910-77)

Grey Walter pioneered robots with a mind of their own. His mechanical tortoises sensed their environments, moving towards light and away from anything they might bump into. They could even find their way back to a charging station when their battery was low. Walter was a pioneer in technologies such as the

electroencephalograph, or EEG machine, for studying the human brain. He claimed these simple robots had the equivalent of two neurons and that by adding more cells they would gain more complex behaviour – something he tried by making a more complex version called Cora (Conditional Reflex Analogue) where he trained the robot to respond to a police whistle in much the same way that Pavlov conditioned dogs to salivate at the sound of a bell. The Cora robot initially made no response to the whistle, but if the whistle was blown when an electric torch was flashed, it soon learned to associate the two stimuli, and responded to the whistle on its own as though it was seeing light.

At around the same time that he constructed his experimental robots, Walter was a member of a very exclusive group of young scientists in the UK known as the Ratio Club. These neurobiologists, engineers, mathematicians and physicists would meet regularly to listen to an invited speaker and then discuss their views on cybernetics, or the science of communications and automatic control systems in both machines and living things. It was one of the very first AI and robot clubs. Most of the members went on to become eminent scientists in their fields. One of the enthusiastic mathematicians was called Alan Turing.



By 1950, Turing had already contributed hugely to the embryonic field of computers. His early work had provided fundamental mathematical proofs, for example that it would not be possible for any computer to predict if it might stop calculating for any given program, or in other words, some problems are not computable. He helped design the very first programmable computers, and his secret work at Bletchley Park helped decode encrypted messages during the Second World War. Like many computing pioneers, Turing also had a fascination for intelligence. What was it? How could you make an artificial intelligence? And if you ever somehow made a computer that could think in the same way that living creatures think, how would you know? Turing decided that we needed a method for measuring whether a machine could think. He called it the 'Imitation Game', but his test became known as the Turing test.

THE TURING TEST

An interrogator can communicate with two people – each in a separate room – by typing text. He can ask any

questions he likes: 'Please write me a poem on the subject of the Forth Bridge'. Or, 'What is 34,957 added to 70,764?' The two people then type their responses. After a while, the interrogator is informed that one of the two people is actually a computer. If he cannot distinguish the computer from the real person, we can then say that the computer has passed the test.

The Turing test became an important measure for AI, but it also drew much criticism. While it may provide some idea of the ability of the AI to reply to written sentences in a seemingly thoughtful manner, it does not measure many other forms of AI, such as prediction and optimization, or applications such as robot control or computer vision.

Turing was not the only pioneer of computers to think about AI. Nearly all of them did. In the US, John von Neumann, a mathematical genius who helped describe how to build the first programmable computers in 1945, worked with Turing on intelligent computers. Von Neumann's last project was on self-replicating machines, an idea that he hoped would enable a machine to perform most of the functions of a human brain and reproduce itself. Sadly, he died of cancer aged fifty-three before he could complete it.

Claude Shannon, another genius who was responsible for creating information theory and cryptography, and who coined the term 'bit' for binary digit, was also deeply involved in the earliest stages of AI. Shannon created a robot mouse that could learn to find its way through a maze, and a computer program that played chess, and in his later years he created other bizarre inventions such as a robot that could juggle balls. In 1955 Shannon and pioneers John McCarthy, Marvin Minsky

and Nathaniel Rochester proposed a summer workshop to gather together scientists and mathematicians for several weeks to discuss AI. The Dartmouth Workshop was held for six weeks in the summer of 1956 and was the first ever focused event to explore (and name) AI. The weeks of discussion resulted in some of the key ideas that were to dominate this new field of research for many decades to follow.

A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence, 31 August 1955

JOHN MCCARTHY, MARVIN L. MINSKY, NATHANIEL ROCHESTER
AND CLAUDE E. SHANNON

We propose that a two-month, ten-man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

The rise and fall of AI

Excitement about AI grew rapidly in the years following the Dartmouth Workshop. New ideas about logic, problem solving, planning, and even simulating neurons fuelled researchers' optimism. Some researchers felt that machine translation would be solved very quickly, because of advances in areas such as information theory and new rules that described how words are put together in sentences within natural languages. Other researchers were investigating how the brain used neurons, connected together as networks, to learn and make predictions. Walter Pitts and Warren McCulloch developed one of the first neural networks; Marvin Minsky designed the SNARC (Stochastic Neural Analog Reinforcement Calculator) – a neural network machine. By the early 1960s even highly experienced and intelligent pioneers were making slightly unrealistic predictions, given the current state of technology.

“ In principle it would be possible to build brains that could reproduce themselves on an assembly line and which would be conscious of their own existence. ”

FRANK ROSENBLATT

AI pioneer in perceptrons (1958)

Fuelled by such excitement, funding blossomed and researchers feverishly worked on machine translation and 'connectionist' (neural network) projects. But the hype was too much. By 1964, funders in the US (the National Research Council) were starting to worry about the lack of progress in machine translation. The Automatic Language Processing Advisory Committee (ALPAC) examined the issues. It

seemed that researchers had underestimated the difficulty of word-sense disambiguation – the fact that the meaning of words depends on their context. The result was that the 1960s AIs made some rather embarrassing errors. Translated from English to Russian and back to English, ‘out of sight, out of mind’ became ‘blind idiot’.

Within our lifetime machines may surpass us in general intelligence.

MARVIN MINSKY (1961)

The ALPAC report concluded that machine translation was worse than human translation, and considerably more expensive. After spending \$20 million, the NRC cut all funding in response to the report, ending machine translation research in the US. At the same time, connectionist research was fading as researchers struggled to make the simple neural networks do anything very useful. The final nail in the coffin for neural networks was the book *Perceptrons* published in 1969 by Marvin Minsky and Seymour Papert, which described many of the limits of the simple neuron model. This spelled the end of neural network research. But it got even worse. Next came the Lighthill Report in 1972, commissioned by parliament to evaluate the progress of AI research in the UK. Mathematician Sir James Lighthill provided a devastating critique: ‘Most workers who entered the field around ten years ago confess that they then felt a degree of naive optimism, which they now recognize as having been misplaced ... achievements from building robots of the more generalized types fail to reach their more grandiose aims.’ The effects of the report had repercussions throughout Europe and the world. DARPA (Defense Advanced Research Projects Agency) cut its AI funding as the agency also realized that promised

results were not being delivered in areas such as speech understanding. In the UK AI funding was discontinued in all but three universities (Essex, Sussex and Edinburgh). AI and intelligent robots had been completely discredited. The first AI winter had arrived.

Despite being deeply out of favour, a few AI researchers continued their work for the next decade. Earlier work was not lost; many advances simply became part of mainstream computer technology. Eventually, by the 1980s there was a new breakthrough in AI: expert systems. These new AI algorithms captured the knowledge of human experts in their rule-based systems and could perform functions such as identifying unknown molecules or diagnosing illnesses. AI languages designed to enable this kind of AI were developed, such as Prolog and LISP, and new specialized computers were built to run these languages efficiently. Soon, expert systems were being adopted in industries around the world and business was booming. Funding was now available again for AI researchers. In Japan, \$850 million was allocated for the Fifth Generation computer project, which aimed to create supercomputers that ran expert system software and perform amazing tasks such as holding conversations and interpreting pictures. By 1985 more than \$1 billion was being spent for in-house AI departments, and DARPA had spent \$100 million on ninety-two projects at sixty institutions. AI was back, and with it came the overexcitement and hype once again.

“ The time at which we might expect to build a computer with the potential to match human intelligence would be around the year 2017. ”

DAVID WALTZ

AI pioneer on reasoning (1988)

But it didn't last. The power of conventional computers quickly overtook the specialized machines and the AI hardware companies went bust. Next it was discovered that the expert systems were horribly difficult to maintain and were prone to serious errors when given faulty inputs. Promised capabilities of AI were not achieved. Industry abandoned this new technology and funding quickly dried up once again. The second AI winter had begun.

The rebirth

Once again, despite being deeply out of favour, some AI research continued. In the 1990s, even the term AI was associated with failure, so it went under other guises: Intelligent Systems, Machine Learning, Modern Heuristics. Advances continued, and successes were simply absorbed into other technologies. Soon, a quiet revolution began, with more advanced fuzzy logics, new, more powerful forms of neural networks, more effective optimizers, and ever more effective methods for machine learning. Robotics technology also started to mature further, especially with new generations of lighter and higher-capacity batteries. Cloud-based computers made it possible to perform massive computation cheaply, and there was so much data being generated every day that the AIs had plenty of examples to learn from. Slowly, but with more and more vigour, AI and robotics returned to the world. Excitement grew yet again, and this time a little fear.

By 2029, computers will have human-level intelligence.

RAY KURZWEIL
inventor and futurist (2017)

“ We should not be confident in our ability to keep a super-intelligent genie locked up in its bottle for ever. ”

NICK BOSTROM

head of Future Of Humanity Institute, Oxford (2015)

By 2019 it was the new AI summer, with thousands of AI start-ups worldwide busily trying to apply AI in new ways. All the major tech companies (Apple, Microsoft, Google, Amazon, Weibo, Huawei, Samsung, Sony, IBM – the list seems endless) were together investing tens of billions of dollars in AI and robotics research. For the first time, AI-based products were being sold to the public: home hubs that recognized voices, phones that recognized fingerprints and faces, cameras that recognized smiles, cars that automated some driving tasks, robot vacuum cleaners that cleaned your home. Behind the scenes, AI was helping us in a hundred tiny ways: medical scanners that diagnosed illnesses, optimizers that scheduled delivery drivers, automated quality-control systems in factories, fraud detection systems that noticed if your pattern of spending changed and stopped your card, and fuzzy-logic rice cookers to make perfect rice. Even if we once again decided not to call it AI in the future, this technology was too pervasive to disappear.

“ Most executives know that artificial intelligence has the power to change almost everything about the way they do business – and could contribute up to \$15.7 trillion to the global economy by 2030. ”

PRICEWATERHOUSECOOPERS (2019)



There had never been so much excitement, so many researchers, so much money, so much hysteria. Despite the ups and downs of AI's popularity, progress in research had never stopped. Today is the culmination of thousands of years of effort poured into some of the most miraculous technologies humans have ever created. If there has ever been a golden age for AI, it is now. Our extraordinary intelligent technology doesn't just help us, it reveals to us the very meaning of intelligence, while posing deep philosophical questions about what we should allow technology to do. Our future is intimately tied to these smart devices, and we must navigate the minefields of hype and misplaced trust, while learning how to accept AI and robots into our lives.

“ Success in creating AI would be the biggest event in human history. Unfortunately, it might also be the last, unless we learn how to avoid the risks. ”

STEPHEN HAWKING (2014)

Each of the following chapters of this book will show you some of the most extraordinary AI inventions so far, and what they might mean for our future. Welcome to the world of AI and robotics.

02 CHOOSE THE RIGHT PATH

'I never guess. It is a shocking habit destructive to the logical faculty.'

ARTHUR CONAN DOYLE

*However of the clock for intuitionists
When of proof, though I not yet, being mine,
Or true a certain portion mathematical
And its element asserted: this was.*

*It was born of the
Axiom-schemas that now I
Not yet fear from implication.*

*One of the concepts of youth, the eye. When
That, there's a mathematics eternal.*

Perhaps this is not the greatest poetry in the world, but this short collection of quatrain, haiku and couplet was generated in a split second by an AI, which was attempting to express ideas about logic with a flavour of the sonnets of Shakespeare. When we read such poems, we

might find some deeper meaning in the words. Somehow the AI captures something that makes us wonder if there is a message being communicated.

In this case, unfortunately, there is not. The poems were generated by a computer following a set of rules that define the structure of each type of poetry. (For example, a Haiku comprises three unrhymed lines of five, seven, and five morae, while a couplet comprises two lines which may be rhymed or unrhymed.) The words were randomly picked from source text (several paragraphs, which included a history of logic, a sonnet from Shakespeare, and a section of text from a 1927 von Neumann paper on logic). Use different source texts and different rules, and the AI will churn out poems about anything you like, in whatever style you have defined.

Symbolic AI

In symbolic processing, words are treated as symbols that relate to each other according to a set of rules. It is almost as though words are objects that you can move around and manipulate, perhaps transforming them, in the same way that the rules of mathematics allow us to manipulate numbers. Symbolic AI allows computers to think using words.

It's perhaps not surprising that symbolic AI was one of the first and most successful forms of AI, because it was built upon the new ideas of logic that had been developed a few decades earlier. Towards the start of the twentieth century, mathematicians such as Bertrand Russell, Kurt Gödel and David Hilbert had been exploring the limits of mathematics to see if it was possible to prove everything, or whether some things that you could express in maths were actually unprovable. They showed us that all of mathematics could be reduced to logic.

'frames' have merged with object-oriented programming languages, and there are many powerful ways that knowledge can be stored; for example, using inheritance, such that a parent object 'tree' might contain the children 'oak' and 'birch', or message passing, such that an object 'seller' might send an argument '10 per cent discount' to trigger a behaviour in another object 'price'. Entire knowledge representation languages have been created, sometimes called ontologies, with their own complicated structures and rules. Many are based on logic, and may be combined with automated reasoning systems to enable them to deduce new facts, which can then be added to their knowledge, or to check the consistency of existing facts. For example, say the AI has learned that 'bicycle' is a type of 'pedal-powered transportation' that makes use of 'two wheels'. If 'tandem' is a type of 'pedal-powered transportation' that also has 'two wheels' then the AI could easily deduce the new fact: 'tandem is a kind of bicycle'. But a different type of 'pedal-powered transportation' with 'zero wheels' doesn't fit the rule, so the AI will conclude that 'pedalo boat' is not a 'bicycle'.

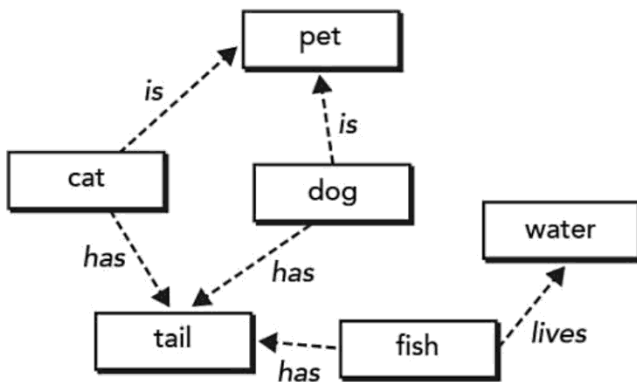
With the growth of the internet, compiling vast collections of facts became easier and easier. Several major efforts in general artificial intelligence aimed at combining enough knowledge together that an AI could start to help us in multiple areas. Cyc was one such effort, which has now been compiling common-sense facts and relationships into a giant knowledge base for several decades.

One computer scientist took this idea even further. William Tunstall-Pedoe created True Knowledge, a vast network of knowledge provided by users on the internet that comprised more than 300 million facts. In 2010 Tunstall-Pedoe decided that since his AI knew so much, he would ask it a question that no human could hope to answer. 'It occurred to us that with over 300 million facts, a big percentage of which tie events,

people and places to points in time, we could uniquely calculate an objective answer to the question “What was the most boring day in history?”

True Knowledge looked at all the days it knew about from the start of the twentieth century and decided that the answer was 11 April 1954. On this day, according to the AI, Belgium held a general election, a successful Turkish academic was born, and the footballer Jack Shufflebotham died. These were fewer events compared to all the other days, and so the AI decided that this was the most boring day. True Knowledge eventually became Evi, an AI that you could speak to and ask questions. In 2012 Evi was acquired by Amazon and became Amazon Echo, the well-known household talking AI.

Symbolic AI is today growing as the internet grows. While AIs such as Cyc and Evi relied on thousands of users providing concepts manually, Sir Tim Berners-Lee, the creator of the World Wide Web, has long pushed the idea that the WWW should become a GGG (giant global graph) of concepts. Instead of just building websites for people to use, the websites should also hold data in a form that computers can understand. Websites were traditionally built like documents, with text, images and videos, or like programs, with behaviours that are triggered when forms are filled and buttons pressed. In Tim’s dream, inside every web page concepts should be labelled with names and unique identifiers. In the Semantic Web, as it is known, websites become databases of concepts, where every element is an object in its own right that can be found independently and has a clear textual label or type. If the entire WWW became a GGG, then our AIs would be able to search, make deductions and reason about the world’s knowledge.



This grand dream for symbolic AI has sadly not been adopted by most web developers in the world, who continue to put vast quantities of data online in a form that AIs struggle to recognize. And the need is becoming urgent. In 2019, it was estimated that 80 per cent of new data will be unstructured – that’s without a consistent representation for knowledge that computers can understand, such as textual documents, images and videos. (Think about all those emails or reports you write as ‘free text’ – that means without explicitly chunking them into labelled sections. Or all the photos and videos you take with your phone – you’re not going through and labelling every scene, or item in frame.) At the same time, the amount of data is growing year by year. By 2019 there were 4.4 billion internet users, an increase of 80 per cent in five years, and 293 billion emails sent daily. There were 40,000 searches per second of the internet with Google and 7,800 tweets per second on Twitter. More and more companies used the internet as part of their business, and generated their own huge quantities of data. In 2016 we were generating 44 billion Gb per day worldwide. It has been estimated that by 2025 we will generate 463 billion Gb per day.

“ I have a dream ... Machines become capable of analysing all the data on the web ... the day-to-day mechanisms of trade, bureaucracy, and our daily lives will be handled by machines talking to machines, leaving humans to provide the inspiration and intuition. ”

TIM BERNERS-LEE (2000)

We no longer have a choice – no human can make sense of these mind-boggling quantities of data. Our only hope is to use AI to help us. Luckily, as we will see in later chapters, other forms of AI are now able to process unstructured and unlabelled data, and tag them with symbolic labels, giving the symbolic AIs what they need in order to think about them. In the end, maybe it doesn't matter if this is true intelligence (strong AI) or simply a kind of 'pretend intelligence' (weak AI). Processing networks of symbols according to rules enables our computers to make sense of our vast universe of data. And perhaps one day, Berners-Lee's dream for the web will come true.

03 WE ALL FALL DOWN

'You don't learn to walk by following rules. You learn by doing, and by falling over.'

RICHARD BRANSON

We see a grainy film footage. A strange wheeled robot looking rather like a wobbly photocopier on wheels with a camera for a head is trundling around a space occupied by large coloured cubes and other simple shapes. In the background we hear the mellow jazzy sound of 'Take Five' played by The Dave Brubeck Quartet. A narration begins, with a high-pitched whine in the background:

At SRI we are experimenting with a mobile robot. We call him Shakey. Our goal is to give Shakey some of the abilities that are associated with intelligence – abilities like planning and learning. Even though the tasks we give Shakey seem quite simple, the programs needed to plan and coordinate his activities are complex. The main purpose of this research is to learn how to design these programs so that robots can be employed in a variety of tasks ranging from space exploration to factory automation.

INDEX

3D printers ref1

2001: A Space Odyssey (Arthur C. Clarke) ref1

2001: A Space Odyssey film ref1

A

Affectiva ref1

affective computing *see* Emotional AI

agent-based modelling (ABM) ref1, ref2

AI (Artificial Intelligence)

in the ref1th century ref2, ref3

artificial neural networks (ANNs) ref1, ref2, ref3, ref4, ref5, ref6, ref7, ref8,
ref9

chatbots ref1, ref2

computer simulations ref1

consciousness ref1

current limitations ref1

emotional/affective computing ref1

general intelligence ref1

human brain simulation ref1

impact on ref1st century society ref2, ref3, ref4, ref5, ref6

optimization ref1, ref2

types of machine learning ref1, ref2, ref3, ref4

see also robotics

AlexNet ref1

American BRAIN Initiative ref1

Ancient Greek mechanical robots ref1, ref2

Andrae, John ref1

Angle, Colin ref1

animals, mechanical ref1

anomaly detection systems ref1

ant-colony optimization ref1

Archytas ref1

artificial immune systems ref1, ref2

artificial life (ALife) ref1

artificial neural networks (ANNs) ref1, ref2, ref3 ref4, ref5, ref6, ref7, ref8, ref9,
ref10

Automatic Language Processing Advisory Committee (ALPAC) ref1

automatons, Ancient Greek ref1, ref2

B

Bayesian AI ref1

behaviour trees ref1, ref2

Berners-Lee, Sir Tim ref1, ref2

bio-inspired optimization ref1, ref2

bodies, evolved robot ref1

boids ref1

Boston Dynamics ref1

Bostrom, Nick ref1

brain damage ref1

brains, evolved electronic ref1, ref2

brains, simulating human ref1

Branson, Richard ref1

Breazeal, Cynthia ref1

Brooks, Rodney ref1, ref2

C