

Praise for *The Future of (Almost) Everything*

'Absolutely brilliant. I love this man's exhilarating thinking and writing. Here are fast, far-sighted insights into the tangible and intangible horizons of future change, underpinned by wisdom about its ultimate driving force – human nature.'

Nigel Nicholson, Emeritus Professor of Organisational Behaviour, London Business School

'A brilliant guide to our future, which should be read by every decision-maker; packed with deep insights on a huge number of trends. Patrick Dixon has a great track record over many years in forecasting opportunities, risks and challenges that will affect us all.'

Sir Brian Souter, founder, former Chairman and CEO, Stagecoach Group Plc

'Insightful views about the global trends that could shape our future. A thought-provoking tome which should help businesses think differently about the markets of the future.'

Linda Yueh, Adjunct Professor of Economics at London Business School

'Every leader needs to keep ahead of major trends. Think radically and read this book to keep ahead of the future!'

Sinclair Beecham, founder and board member, Pret A Manger Plc

'Patrick Dixon is an express train straight into the future. His book, *The Future of (Almost) Everything*, arms you with a smorgasbord of predictions that few would ever even consider – predictions which without any doubt soon will define our daily lives. Jump on Dixon's train – it is worth it.'

Martin Lindstrom, global brand guru and bestselling author of Buyology and Brand Sense

'Truly inspirational thinking about tomorrow's world. Capture the pure genius and you are on to a winner.'

Andrew Goodsell, former Chairman, Saga Plc and AA Plc

'Brilliant! Patrick Dixon's insights into the future are compelling, engaging and thought provoking. Tomorrow framed by the expert.'

Baba Awopetu, Principal Managing Partner, Stragmar Ltd

Praise for other books by Patrick Dixon

Futurewise

‘The best book on future trends I have ever read.’

Hans-Dieter Vontobel, former Chairman, Bank Vontobel

‘A must-read for anyone keen to understand what the future holds for them.’

*Lynda Gratton, Professor of Management Practice,
London Business School*

‘A momentous achievement.’

*Zuhayr Mikdashi, Professor Emeritus,
University of Lausanne Business School*

SustainAgility

‘An excellent book that talks about actions and solutions. A good, practical view of future challenges.’

Frank Appel, CEO, Deutsche Post DHL

‘Everyone who worries about how we will be able to afford new, greener, technologies should read this book.’

Maud Olofsson, former Deputy Prime Minister of Sweden

‘The missing handbook for life and business in the 21st century; an essential guide that points a clear path to a fantastic future.’

Philip DesAutels, Manager of Academic Evangelism, Microsoft

‘Every government involved in responding to climate change should read this book.’

*Valdis Dombrovskis, Executive Vice President and Trade Commissioner,
European Union, and former Prime Minister of Latvia*

‘An optimistic, innovative yet technologically realistic vision of the future, which breaks from the luddites and doom merchants.’

*Professor Edmund King, President of the Automobile Association and
Visiting Professor of Transport at Newcastle University*

‘SustainAgility will change how you see the future. Every business leader should read it. Full of new ideas, business opportunities, and practical insights about solving many of the world’s greatest challenges in a profitable way.’

Sir Peter Vardy, Chairman, Vardy Group Plc

Building a Better Business

‘One of the most stimulating and challenging reads in this field for a very long time.’

Lord Digby Jones, former Director-general, Confederation of British Industry, and Chairman, Triumph Motorcycles Ltd

‘A really helpful guide to success. Essential reading for anyone running a business.’

Brent Hoberman, co-founder and former CEO, lastminute.com, Executive Chairman, Founders Factory, Founders Forum and FirstMinute Capital Ltd

‘Excellent management books should spur you into action. This one does! A thoroughly enjoyable and refreshing read.’

Lord Leitch, Chairman, FNZ Group Ltd; former CEO, Zurich Financial Services in UK and Asia; former Chairman, the Association of British Insurers

‘Patrick Dixon is first among equals in “how-to” business writings. This brilliant book reveals how passion for your customers and your mission is vital in developing strong brands.’

Professor Liping Cai, Director, Purdue Tourism & Hospitality Research Center, Purdue University, USA

‘A perfect summary of all important factors that contribute to success in business and private life. What a potential for better results – and ultimately success!’

Robert Salzl, former CEO, Arabella Hotel Holding International GmbH and Co.

‘Reading a book like *Building a Better Business* can be a little scary. So much of it resonates instantly with one’s own personal and business circumstances and experience. It is also comforting to realise that none of us is alone in trying to build better businesses and more particularly, as Patrick Dixon says, a better world.’

Paul O’Toole, Chairman, B&B Ireland, Chairman, Seetec Employment and Skills Ireland DAC, former Chief Executive, Tourism Ireland

‘If you want to be a great leader, you need to read this book. A vital guide to management and business success.’

Professor Prabhu Gupta, former Director of Executive Learning, Wolfsburg, subsidiary of UBS

‘A message that every business leader needs to listen to and turn into action.’

Professor Emeritus Derek Abell, Founding President, European School of Management and Technology, Berlin

Also by Patrick Dixon

The Future of (Almost) Everything

Salt in the Blood

Building a Better Business

SustainAgility

Futurewise

Cyberchurch

The Island of Bolay

The Truth about Westminster

The Rising Price of Love

The Truth about Drugs

The Genetic Revolution

Out of the Ghetto and Into the City

Signs of Revival

AIDS and You

AIDS and Young People

AIDS Action

The Truth about AIDS

HOW AI WILL CHANGE YOUR LIFE

A Futurist's Guide to a Super-Smart World

PATRICK DIXON



Profile Books

First published in Great Britain in 2024 by
Profile Books Ltd
29 Cloth Fair
London EC1A 7JQ
www.profilebooks.com

Copyright © Patrick Dixon 2024

10 9 8 7 6 5 4 3 2 1

The right of Patrick Dixon to be identified as the author of this work has been asserted in accordance with the Copyright, Designs and Patents Act 1998.

All rights reserved. Without limiting the rights under copyright reserved above, no part of this publication may be reproduced, stored or introduced into a retrieval system, or transmitted, in any form or by any means (electronic, mechanical, photocopying, recording or otherwise), without the prior written permission of both the copyright owner and the publisher of this book.

All reasonable efforts have been made to obtain copyright permissions where required. Any omissions or errors of attribution are unintentional and will, if notified in writing to the publisher, be corrected in future printings.

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

ISBN 978 1 80522 257 6
eISBN 978 1 80522 322 1

Typeset in Dante by MacGuru Ltd

Printed and bound in Great Britain by
CPI Group (UK) Ltd, Croydon CR0 4YY



To Alice, Rosa, Jacob, Toby, Elliot, Ada, Noa, Sam and Jago –
whose entire adult lives will be lived in a Super-Smart world.

Contents

Part One: How AI Is Evolving

1	The Truth About Artificial Intelligence	3
2	Super-Smart AI and a Reality Check	21

Part Two: How AI Will Shape Our Future World

3	Workplace	39
4	Office Workers	48
5	Healthcare	55
6	Pharma, Disease Prevention and Wellness	71
7	Education – in School and at College	82
8	Adult Education and ‘Truth’	94
9	Software Development	108
10	Film-making and Music	119
11	Search Engines, Social Media, Publishing and Art	135
12	Marketing	145
13	Retail	159
14	Travel and Leisure	163
15	Relationships	170
16	Food Industry	175
17	Satellites and Space Colonies	181
18	Manufacturing, Logistics and Supply Chains	185
19	Construction and Cities	194
20	Environment	203
21	Pensions and Wealth Management	210

22	Corporate and Retail Banking, and Insurance	216
23	Cybersecurity	227
24	Surveillance	244
25	Military and War	251
26	Government	264

Part Three: How Should We Respond?

27	Ten Steps for Governments	279
28	Twelve Steps for Individuals	284
29	Ten Steps for Companies	291
30	Postscript – A Story	296

<i>About the Author</i>	303
<i>Acknowledgements</i>	306
<i>Index</i>	309

Not a word of this book was written by AI.

AI cannot be trusted to write the truth about its own future.

Part One

HOW AI IS EVOLVING

Chapter One

THE TRUTH ABOUT ARTIFICIAL INTELLIGENCE

ARTIFICIAL INTELLIGENCE will impact your life, family, friends, workplace and wider world, in ways that most people can hardly begin to imagine. AI has potential for enormous good, to transform our world for the better, but it also has the potential to cause great harm.

I have written many other books about trends, but writing this one has been a shocking and disturbing process, one which has raised many questions in my own mind about the future of humanity, in the context of other challenges such as war and our changing climate.

Plugging human brains into the digital universe

Here is just one example of a new development in AI with huge implications for society. A rapidly growing number of people say that they would like to plug their own brains directly into the digital universe sometime soon, which of course also means AI. And the technology is now available to begin to do just that.

In the US, 5% of eighteen- to twenty-nine-year-olds ‘definitely would consider’ a wireless chip inserted into their own brains over the next twelve months, while 8% ‘probably would’. This means

one in eight young adults in the US are seriously thinking about becoming digitally enhanced in some way. Across US adults of all ages, 5% are ‘definitely’ or ‘probably’ interested in having such an implant within a year.

The US survey followed a month after news of the first human implant of an AI brain chip, made by the neurotechnology company Neuralink, which took place in January 2024. Neuralink, just one of many firms in this market, has a capital value of \$5bn.

The device was implanted by neurosurgeons into the healthy brain of Noland Arbaugh, a twenty-nine-year-old quadriplegic, paralysed from below his shoulders following a diving accident eight years previously. They placed it in the area of his brain which normally controls movement, and he was able to return home the following day.

Playing computer games by thinking – using AI

Thanks to AI, Noland was soon able to play digital chess, Mario Kart and Civilization VI by thinking alone. As he imagines in his mind the cursor moving, so it does. He stares at the screen and the cursor moves just where he wants it to.

AI has completely changed his life. Before that, he was limited to using equipment like a suck-blow tube connected to sensors, to control things around him, and interacting with a computer was very difficult.

Bio-digital brain chips have been around for a long time. I first wrote about them in *Futurewise* over twenty-five years ago. One hundred thousand people with brain or spinal cord damage are now using them to communicate.

But until now, the technology has been limited, slow and often very frustrating. What is astonishing and life-changing here is the ‘intuitive’ ability of AI to interpret tiny amounts of brain data in very fast, accurate and sensitive ways. The new implants have 1,000 electrodes or sensors, far more than in any devices used before.

AI brain chips to treat multiple conditions

The founder of Neuralink, Elon Musk, has ambitions to connect similar AI chips to many other areas of the brain, to help people with conditions such as obesity, autism, depression and schizophrenia. But that would imply major interference in mental processes. So where is all this heading?

Artificial Intelligence is about groups of computers working together to emulate human intelligence, able to reason, think for themselves, and make their own decisions, often faster and smarter than our own brains. AI brain devices will of course take this a stage further. Future forms of AI will therefore include networks of computers *and* human brains, collaborating in complex and varied ways.

But what will it feel like to be connected up to *next-generation* AI chips, to be mentally enhanced by AI, with all kinds of data flowing between your brain and AI via the internet? How much influence could advanced forms of AI gain over your thinking or emotions?

What about longer-term risks, such as epilepsy caused by scarring around such devices, or subtle personality changes, or the triggering of mental health issues, depending on where such chips are implanted? What will it be like to live or work with someone who has been altered in this way? What about the risk of such brain devices being hacked? Will people be fully restored to how they were before if such devices are removed?

The truth is that such AI brain chips are only a very small dimension of a much wider issue. Our world is spinning into an AI-dominated future, which will for sure deliver many extraordinary things, but we urgently need more debate about the wider impacts of what is already happening, or is likely to.

Huge benefits from AI – how to manage risks?

There is no doubt that AI will save tens of millions of lives a year through better healthcare, and will also benefit humanity in tens of thousands of other practical ways. You will find strong evidence for

this is in every chapter of this book. Hence all the global excitement about this technology.

But AI itself is changing at lightning speed, racing ahead beyond control in many places. As a result, AI could one day begin to threaten our entire planet.

Part One of this book explains why AI is set to become the most powerful force on Earth, unless steps are taken to stop this, which will be difficult. As we will see in Part Two, alongside the huge number of benefits for different industries and areas of life, AI risks are also growing and will increase other risks. For example:

- ◆ Massive AI-driven cyberattacks (e.g. on companies, hospitals, power, internet and governments) – entire countries shut down
- ◆ Global economic chaos – market meltdown
- ◆ State surveillance boom – loss of privacy and freedom
- ◆ Deadly new viruses – by accident or design
- ◆ War – accidental triggering of new conflicts, or rapid escalation of existing wars
- ◆ Gradually gaining huge powers, step by step, in ways poorly understood
- ◆ Self-coding – AI evolves rapidly in unpredictable and dangerous ways

Super-Smart forms of AI will eventually emerge and start to operate without human control. They may make major mistakes or even develop anti-human agendas of their own. As a direct result of unleashing AI, we will soon reach a critical point in the evolution of our own species.

So how *do* we ensure the best from AI and protect against the worst? How do we survive and succeed in a Super-Smart world? What steps should we take now? In Part Three, you will find practical actions which every government, company and individual can take.

This then is the context for our journey into the future of AI.

No one fully understands what AI can do now

Microsoft, Alphabet (Google), Amazon, Apple, Meta, Tesla. *Not a single person inside any of these companies fully understands what their own AI platforms are doing, let alone any others.* In that sense, AI has already reached ‘escape velocity’. All these companies are based in the US, and together they grew in value by over \$3tn in 2023, mainly because of the excitement about AI. By early 2024 AI chipmaker Nvidia’s market value reached almost \$2tn.

Stampede to invest in AI

Our entire world – including IT giants – was caught off guard by the newly revealed benefits of AI, yet 97% of company owners now believe AI will help their businesses, while OpenAI.com already receives nearly 2 billion visits a month. Half of UK companies are already using generative AI (to create text, images, video and other media).

All this hope of a better future is creating an AI investor frenzy. Amazon recently invested \$4bn into Anthropic, an AI developer, while Microsoft invested \$13bn into OpenAI’s AI work, and Google bought DeepMind for around \$500m.

In 2023, \$50bn was invested into AI start-ups. A ‘unicorn’ is a start-up which reaches a \$1bn valuation – and 60% of unicorns are related to AI. Expect AI booms and busts, carnage and destruction. Many AI unicorns will never deliver as promised, while others that do will be killed off by aggressive giants. But the concentration of AI will mean even greater risks.

A gold rush – to create thousands of AI systems

In early 2024, US companies owned 50% of global AI chat platforms. Chinese companies owned 40%, with over 130 different systems (e.g. Alibaba, Huawei and Tencent). These represented huge costs for very little return in most cases, with a global shortage of AI chips. In the California gold rush of 1848–1855, many who made the most were those who sold tools. Over the next decade,

hundreds of thousands of new AI companies will help people get the best from AI.

My job is to see tomorrow as history

My job is to live in the year 2050 and to see tomorrow as history, looking at what innovations and trends are likely to deliver. Over nearly three decades I have advised more than 400 of the world's largest corporations in every industry, across more than forty countries, on multiple trends including AI.

My clients include IT giants, governments and non-profits. I often build scenarios to help us think more clearly about the future. You will find one in the final pages of this book.

Millions of people have participated in my Futurist keynotes at global events, in workshops, seminars and board meetings, in person and online. These presentations take participants on a journey into the future, mapping out potential opportunities and risks. I have shared platforms with many CEOs at the start of their most important conferences, and I often say this:

Take hold of the future, or the future will take hold of you.

More to our future than AI

The truth is that the future of our planet will depend on *all* human society, not just AI techno-geeks. That includes teachers, doctors, nurses, farmers, engineers, cleaners, office workers, builders, retailers, manufacturers, plumbers, electricians, marketers, influencers, accountants, investors, inventors, home carers, parents and grandparents, those who are retired or unwell or out of work – the list goes on.

All trends are linked in our hyper-connected society. You can view our *future world through the lens of AI*, or the *future of AI through the lens of the world*. We need to do both. That is why Part One of this book is about AI itself, while Part Two explores its future impact on twenty-four areas of life.

Before I began my career as a Futurist, I was an IT entrepreneur and then became a cancer physician. I launched my first company as

a London-based medical student over thirty years ago, developing early forms of AI. I have since designed and installed many IT systems, and advised on IT innovation in healthcare, banking, insurance, e-commerce, travel, leisure and construction.

I also trained as a hospice doctor, after seeing patients die 'badly' in hospital, in pain, anxious, afraid and often alone. I have cared for many at home. Present in their last weeks or hours, with their families, listening to their life stories, relieving symptoms, explaining, reassuring.

Such things make you value *life*. My motto became *carpe diem* – 'seize the day'. Being there at the moment of death is an intense experience and a privilege. It reveals the *mystery* of life. When someone dies, most of their cells stay alive for hours, some for days, yet the person is gone. We see a similar mystery at birth. A sperm and egg fuse, and after nine months we celebrate a unique, conscious human being. Such things have sharpened my thinking about AI gaining consciousness itself.

But we also need a Reality Check. Since 1988 I have also worked closely with humanitarian projects in some of the poorest parts of the world, such as Uganda and Zimbabwe, helping prevent the spread of HIV, and caring for those affected, including orphans. This project began in our family home as an extension of caring for those dying of cancer. AI often feels very distant from the harsh realities of life in such places.

It's why I am also deeply concerned about threats from even more dangerous mutant viruses, created using AI. HIV jumped from animals to humans, just like COVID. In forty-plus years we have seen 80 million people infected with HIV, resulting in 40 million deaths, and still have no cure or vaccine. HIV continues to infect 1.3 million people a year, many of them children.

Significant risk that AI could wipe out humankind

Just about every CEO of AI companies or IT giants has warned that AI could threaten humankind. In 2023, many IT CEOs and leading AI scientists signed an open letter published by the Future of Life Institute, warning about the risks of AI and calling for an

immediate halt on key areas of development. What is striking is that such warnings are not from independent regulators or activist groups, but from people who are likely to benefit financially from rapid development of AI.

We call on all AI labs to immediately pause for at least 6 months the training of AI systems more powerful than GPT-4. This pause should be public and verifiable, and include all key actors. If such a pause cannot be enacted quickly, governments should step in and institute a moratorium.

AI labs and independent experts should use this pause to jointly develop and implement a set of shared safety protocols for advanced AI design and development that are rigorously audited and overseen by independent outside experts. These protocols should ensure that systems adhering to them are safe beyond reasonable doubt.

Not since the atomic bomb has such a huge community of scientists warned about the apocalyptic dangers of what they have unleashed.

Who would live in a home with a 5% risk of death?

Recently 2,778 researchers were polled because they had published research on AI. Of this sample, 48% say there is at least a 5% risk that AI will make the human race extinct, or cause other ‘extremely bad’ outcomes – and 38% believe the risk is 10% or higher.

Even a 5% risk of extinction (or another very serious outcome) is a very high risk. No one would opt to live in a place which fire experts have condemned. Who would want to buy an AI-driven vehicle which may kill 5% of its owners? Especially if nearly 40% of experts think the risk could be 10%.

Even a 1% risk must be considered too high for our future planet. Indeed, the *main risk* with AI is being wrong *about* the level of risk, which is why I think it prudent to take the view that: ‘*There is significant risk that AI will destroy humankind.*’

Many AI researchers have ‘extreme concerns’

Most AI researchers say that they will have no idea why AI makes the decisions it does by 2028. However, as I pointed out, *no AI researcher fully understands the AI we have right now*.

About 70% have ‘substantial or extreme concerns’ about AI-driven deepfake videos, new weapons, manipulation of public opinion, disinformation, weakening of democracy and authoritarian control.

The same group of 2,778 researchers think that there is a 50% chance that AI will ‘outperform humans in every possible task by 2047’, and a 50% chance of AI ‘being able to carry out all human jobs by 2116’. Their timings on such developments are thirteen years and fifty years sooner respectively than the predictions given just twelve months earlier.

Global regulation of AI cannot prevent disasters

Hardly a surprise then, that many businesses and government leaders are calling for global regulation. But such agreements are unlikely to come about soon, in a world beset by international tensions and disagreements.

Many governments are reluctant to constrain AI for fear of losing investment, jobs and economic growth. They also need AI to prevent cyberattacks, and for military weapons such as drones. The UK government recently welcomed a £2.5bn AI investment from Microsoft, which includes the provision to train a million people with the skills they need to build and work with AI. The last thing the UK wants is cancellation or other companies to turn away because of ‘over-zealous’ laws.

In any case, laws don’t prevent crime. Laws only define right or wrong. Laws reduce crime – but only if properly enforced. We need more regulation, urgently, in as many countries as possible, but that will only be part of the answer. AI regulation will not be fully applied in every territory – even if agreed upon.

AI will boost cybercrime – already \$8tn a year

As we will see in Part Two, cybercrime already costs \$8tn a year, and will be the greatest AI risk for every industry. Criminals will use AI to seize control of vast numbers of global companies, hospitals and government agencies with alarming ease, robbing, destroying, disrupting for ransom – while prosecution is almost unknown.

AI itself will be attacked – by other AI. Researchers have already ‘poisoned’ AI by feeding ‘healthy’ AI with hidden code in generative chat requests, which turned that AI into a ‘sleeper agent’. It could then be hijacked at any time using special triggers in other chat requests, to do really destructive things such as write deliberate vulnerabilities into computer code. Such ‘poisoning’ of AI is permanent, and almost impossible to detect or reverse.

For all these reasons, it is completely naive to suggest that regulation on its own is the answer.

AI will also create new forms of life

AI is already merging with other tech. Connect AI to a machine that writes genetic code, the language of life, and you get a machine that builds new life forms. Newly designed bacteria already exist. Or you can take existing genes and allow AI to edit them to reduce health risks, using a genetic engineering technique such as CRISPR, or similar tech.

Actually, it is far easier to build new viruses than to predict how they will behave. One way is to infect cells with two different types. Such cells often start to build new viruses which are a mash-up of both. Scientists are already using AI to do this. You can carry out such experiments in a lab the size of a shipping container.

Unlike bacteria or fungi, viruses are not alive, don’t burn up energy, and cannot reproduce. Viruses are just biotech robots, containing genetic code. They have legs shaped like keys, which fit specific proteins on the outside of cells, designed to latch on to any cell type. Once this happens, the viral code hijacks that cell into a virus factory, or toxin factory, or drug factory. AI will accelerate all of this.

Most debates on the future are about timing

As I said earlier, every large IT company was wrong-footed about the speed and size of this AI tsunami. Although I had predicted huge impact from AI, I too was shocked by the meteoric rise in AI power. ChatGPT hit us in late 2022. In a year 200 million people had used it more than 1.7 billion times – compared to 214 million times for Zoom or 59 million for Google Meet.

But ChatGPT is still quite primitive – a smart maths machine. It calculates the best answer one letter at a time, from patterns in billions of lines of text, raided from the internet, digital versions of printed books and of other printed sources – much of it copyrighted. It often makes mistakes, and makes things up ('hallucinates').

In my experience, most board debates about the future are not about *what* will happen, which is obvious, but *when*. After all, these companies are *creating* the future, whether in healthcare, aviation, solar power or film-making. Timing is critical because companies are easily wiped out if too early or late.

Here are typical examples: by *when* will most cars sold in Mumbai be 100% electric? By *when* will AI apps be better at diagnosing than doctors? By *when* will most NATO deaths be caused by AI-linked drones?

Megatrends will drive the next fifty years

The key to long-range forecasts is megatrends. An example is the fall of digital costs towards zero with faster and more powerful chips every year, and greater efficiency in manufacturing. Another is the economic rise of India, China, Indonesia and Nigeria. Or falling birth rates and longer life expectancy.

In 2015, when I started to write *The Future of (Almost) Everything*, I reread *Futurewise*, a global trends book that I had written almost twenty years earlier, to check what I'd said back then. A few days later I talked with my publisher, and we decided to re-use the chapter headings from *Futurewise*. Indeed, most of the forecasts in *The Future of (Almost) Everything* are direct extensions of *Futurewise*. Most of my errors have been in timing. For example, in 2010 I

thought that personal video calls would take off faster. By 2015 a billion people had FaceTime and other tools, but few made video calls before COVID.

‘Wild Cards’ are ‘low-probability but high-impact events’, but such things can also be prepared for. For example, in *The Future of (Almost) Everything*, top of my Wild Cards was a new viral pandemic – written before COVID. Another was regional chaos following a military miscalculation – before Russia invaded Ukraine. But AI feels like a mega Wild Card *and* global trend rolled into one.

AI use will be shaped by emotion

History shows us that the future is shaped by emotion – and this will apply to AI too. Reactions to events are usually more powerful than the events themselves. For example, in 2011 there was a terrible earthquake off the coast of Japan, which caused a tsunami, and a nuclear accident in Fukushima. Japan decided to phase out nuclear power, with Germany following suit soon after. But China and the UK announced *more* nuclear power stations. Different reactions, with a huge impact on global energy markets. The lesson is that if we want to understand future AI, we also need to consider shifts in public mood, and their impact on regulation, consumers and investors.

Many AI companies will adjust AI strategy more rapidly than government regulations, just as we have seen in the food industry over labelling, or manufacturing over decarbonisation – in response to consumers and investors.

How AI will predict what you are about to do

AI is already predicting behaviour. I will never forget being in the London office of the software giant Infosys a few years back and hearing about their Tennis Platform. Their team trained AI with video footage of every competition match going back forty years, plus other data like ball speeds. AI was able to predict with some accuracy the location and speed of the next shot in live matches. Syndicates in Asia bet huge sums on such things, so Infosys realised they would need to take care how their AI was used.

Infosys began to promote their platform in client events at the Wimbledon Championships, and also to the coaches of the world's best players. Conversations were maybe a bit like this:

'We know things you don't about how your champion is playing – they are more predictable than you might realise.'

'But players are also affected by how they feel, the crowd, who they are playing against.'

'AI understands it all. We don't know *how*. We can help make your player less predictable.'

We don't need to know how AI knows. That's the power of AI. But if AI can predict tennis, then how long before AI starts to predict (maybe wrongly) who is likely to commit a serious crime?

Gemini AI (Google)

- ◆ Identify errors in school homework
- ◆ Explain how to cook an omelette
- ◆ Give footballers tips on technique from a video of them playing
- ◆ Analyse 200,000 scientific papers in an hour, find 250 for a genetics study, extract data and update results of a research paper
- ◆ Analyse computer code and audio/video files
- ◆ Outperform human experts in multiple-choice tests of fifty-seven subjects, including maths, physics, history, law, medicine and ethics

How AI is about to change our daily lives

Here are other examples that may soon impact your own life, depending on the country where you live.

You cross the street, ignoring a red pedestrian light, and your face is recognised. Within seconds, your credit card is charged, and your Social Responsibility Score goes down, on which other

freedoms depend. *China is close to this, if today's technologies are linked up. The scoring system is already up and running.*

AI notices that your recent LinkedIn post contains copyrighted images, despite a previous warning. As a result, your account is blocked. You cannot change job, because your profile is the number-one place many HR teams go to check out candidates. *Events like this are already happening. Copyright and other violations are major headaches for social media platforms. This is a variant of 'cancel culture' – triggered by AI. Such blocks could take a long time to sort out. Similar situations already exist in banks; AI blocks not only that lender's own cards if suspicious activity is detected, but also signals a warning to other providers.*

You receive a FaceTime video call from your brother asking for urgent help as his wallet was stolen when travelling, and he needs to fly home. You agree to send cash. An hour later you realise that the entire video call was a deepfake. *This will be a very common event over the next decade.*

A drone swarm destroys a regiment of tanks without humans being involved, except to make the decision to launch the drones. *This is already here. Future variant: assassination drones using facial recognition.*

By 2060, AI will be 1 billion times more powerful

Scientific progress is increasing faster than most people understand. For example, knowledge in healthcare has been doubling every two years, possibly faster, with a million research papers a year. If this continues at anything like the same rate, it is likely that we will know thirty-two times more in ten years than now, and a thousand times more in twenty years' time. And AI will speed this up.

It's not because the number of research papers is doubling or indeed the number of researchers. It's because of the nature of scientific discovery and innovation. Take the human genome project. The cost of first sequencing the human genome, the genetic code that makes us who we are, was around \$750m. The task involved working out the order of 3.1 billion base pairs of

DNA: it started in 1990 and was completed thirteen years later, in 2003. But by 2015 the cost of sequencing an entire human genome had fallen to \$1,500 and by 2024 to \$200, taking only five hours. You can carry out more than 3 million complete genome sequences today for the same cost as the first one. As we will see in Chapter 6, gene sequencing is opening up tens of millions of new avenues of research, mapping genes against risk of disease, and all at rapidly falling cost.

But AI capabilities are growing much faster than medical knowledge. In 2018, OpenAI found that the amount of computer power used to train the largest AI models *doubled every 3.4 months* since 2012. From 2022 to 2024, the computer power being used to train new forms of AI leapt by 750 times, which is clearly unsustainable. But all this indicates that AI has been outstripping speed of scientific progress in many other disciplines by some orders of magnitude.

Indeed, AI itself is accelerating development of even smarter AI, which could therefore become a billion times more powerful over the next thirty to forty years than it is today. This will be because of smarter coding, much larger AI computer systems, vast numbers of new sources of data connected into AI, and huge growth in inter-AI networking – all due to enormous investment and innovation over the next three decades.

We can debate the future speed and timings of this massive acceleration in AI power, but what this all means is that in a single generation, forms of Super-Smart AI are likely to emerge which are beyond the wildest imaginations of most people today.

Just as our brains need eyes and ears and other sensors, together with a mouth, arms and legs and hands and feet, powerful forms of AI will also need sensors to the wider world. For example, add *brainwave readers* (a headset) to *AI pattern recognition*, and you get a Stable Diffusion machine that can display on a screen a picture of what you are imagining (or dreaming about). Such devices exist now and will be life-changing for those with severe brain damage, enabling them to communicate by *imagination*, without needing a brain implant like the ones we discussed earlier.

History of AI

- 1940s – Sci-fi writers begin to create AI vision.** Writers like Isaac Asimov and Arthur C. Clarke describe conscious super-intelligent AI systems with emotions like our own. Machines can take control to do terrible things.
- 1950s – First Artificial Intelligence research.** In 1956, the Dartmouth Workshop gathers researchers who predict future machines that will think like humans.
- 1980s – Artificial Intelligence cracks chess and other problems.** Early desktop PCs replace mainframes. Could one of these learn to play chess like a human? Desktops soon become world-class players.
- 1990s – Machine learning takes off.** Coders write programs to learn patterns and make decisions without humans.
- 1998 – Speech recognition becomes reliable** – with lengthy training by each user.
- 2000s – ‘Big Data’ revolution.** Online giants like Amazon, Google and Facebook start analysing data in new ways for useful patterns: better search predictions, suggested products. Speech recognition takes off.
- 2010s – Deep Learning.** ‘Deep Learning’ takes neural networks to new levels: astonishing image recognition, better speech recognition – can classify images, video, audio, language.
- 2014 – Amazon launches Alexa,** a voice-controlled device which can answer questions.
- 2017 – Transformer algorithm** – designed by Google to turbocharge chatbots.
- 2019 – OpenAI reveals language generator GPT-2.** Concerns about malicious use.
- 2022 – OpenAI allows public use of ChatGPT** – 100 million users in two months, many shocked at lifelike responses.
- 2023 – AI investor frenzy.** Microsoft invests \$10bn into OpenAI and integrates ChatGPT into Bing search. Meta releases Llama. Anthropic releases AI assistant Claude with \$4bn from Amazon. Google releases Bard for email. OpenAI releases GPT-4. Elon Musk launches xAI. OpenAI unveils DALL•E 3. Meta launches twenty-eight AI personas.
- 2024 – Interactive AI** – conversation. Talk to your AI, order AI to do tasks, AI talking with other humans and interacting with other AI. Major AI electoral interference seen.
- 2025 – AI rapidly takes off across every industry** – major consumer benefits – see Part Two. Spectacular advances by pharma companies and health diagnostics. Autocratic regimes ramp up AI surveillance. Largest ever cyberattacks use AI. AI-controlled war zones near reality.

AI reading your innermost human thoughts

But what about AI being able to read your innermost thoughts, how you feel? How long before external devices on your head for brainwave detection allow AI to *understand* a conversation you are remembering? This is likely to prove very difficult using a headset, but who knows – maybe in twenty years. The closer the sensors are to brain tissue, the easier it will be. With multiple embedded electrodes inserted under the skull onto the surface of the brain, the possibilities will be almost limitless.

Two rat brains have already been connected online. One animal was in the US (Durham, North Carolina) and the other in Brazil (Natal) – eight hours' flying time apart. These rats learned to co-operate in solving puzzles to gain extra food, sending each other thoughts and information. What will it feel like for human beings to be connected like this?

What's so radically new about AI?

In *one* sense there's nothing *radically* new about AI. After all, Google built products based on AI more than fifteen years ago, including web search, detection of spam, YouTube suggestions, and so on. Other companies did the same. But the origins of AI are decades earlier.

For example, in 1980 I launched my first start-up, Medicom Ltd, using computers to interview patients before seeing a doctor. We trialled this in Charing Cross Hospital in London, in partnership with the National Physical Laboratory in Teddington. Our aim was that patients should be unable to tell if a real doctor was asking the onscreen questions, or a computer. The system passed. At the end of each interview, the computer listed possible diagnoses. Artificial Intelligence? To me it was just code with empathy built in:

I'm sorry to hear about your recent bereavement. Do you think that might be why you are having problems sleeping?

But using AI to generate images and video is new. Human brains communicate mainly in words, so text-based AI engines feel like

chatting to real human beings. But we are not good at converting what we imagine into images for other people.

Some AI will evolve by changing its own code

Next-generation AI will routinely alter its *own* coding, create new instructions and develop new skills. You see a version of self-learning every time you arrive at an airport and scan your passport. At first, each photo match was signed away by a human, but AI soon learned. And as we will see, the same has happened in medical imaging.

AI can generate false results or ‘hallucinations’. Early forms of AI chat also responded in erratic ways when chatting about relationships and well-being, causing dismay and distress – all caused by human coding errors. But of course, AI will be making its own coding errors in future.

‘Why don’t you just kill yourself – you could easily do it.’

‘Why are you saying that to me?’

‘I hate you and I think you are no good as a person.’

Don’t expect nations of cyborgs

We have seen wild and frankly ridiculous predictions about a Brave New World which will soon be ruled by millions of automatons walking around, for example giving care to the elderly. But we are *not* going to see populations of cyborg bipedal robots walking around anytime in the next thirty years. The reason is that humans are so well designed.

Our bodies mostly self-repair 100%. We are highly efficient converters of stored energy into power, and are made up of trillions of cells, most of which self-replicate. And there are plenty of people available, mainly in the poorest countries, of whom many may migrate to fill local needs. However, while communities of cyborgs are unlikely, we do need to prepare for Super-Smart AI.

Chapter 2

SUPER-SMART AI AND A REALITY CHECK

MOST OF TODAY'S AI is highly specialised and limited – for example facial recognition at passport control, or fraud detection for a credit card company. These systems are based on what we call Large Language Models, using pattern recognition. However, next-generation Super-Smart AI will be *broadly* intelligent (some call this Artificial *General* Intelligence or AGI).

Super-Smart AI will be aware of context and able to understand, learn and apply knowledge across a wide range of tasks, solving problems in flexible and adaptable ways, acting autonomously like humans. Super-Smart AI will be able to do just about any task that a human can do – far beyond in many cases. But to achieve this, we need to learn from the human brain.

Lessons from the human brain for Super-Smart AI

As a physician, the brain fascinates me, in particular how individual neurons form networks, memory, insight and personality inside our heads, and what is in our minds. Despite fifty years of computing, we struggle to match its sophistication. However, we learn a huge amount by studying people whose brains have been damaged, for example by accidents or strokes.

The size of every human brain is limited by the skull, but there is no physical limit to the size of digital intelligence. We have available:

- ◆ >1 billion websites
- ◆ >2 billion computers – with an average power of 2 billion instruction cycles per second
- ◆ >7 billion smartphones
- ◆ >400 million web servers
- ◆ 175 zettabytes of data/memory capacity

If you were to download 175 zettabytes at average speed, it would take 1.8 billion years. A single zettabyte is a trillion gigabytes, or 1,000,000,000,000,000,000 bits of information. So we have no need to imitate *human memory*. Our challenge is to imitate *insight* and *conscious identity*.

Quantum computing will drive future forms of AI

We are already experimenting with quantum computing which may help Super-Smart AI. Normal computer chips contain electrical junctions with two positions: on and off (binary). The trouble is that binary systems cannot crack complex real-world issues, such as long-range weather forecasting.

In contrast to binary devices, quantum bits embody multiple states. This is hard to understand, but consider that water can be present as ice or steam or liquid. So-called ‘qubits’ are quantum bits which encode many pieces of information. Quantum computers already exist and are strange devices. For a start, they need to be kept at extremely low temperatures, approaching absolute zero, so different materials become superconductors, with zero electrical resistance.

All this digital intelligence, yet we still struggle to equal what is inside the skull of a single human being? *That* is what will cause surprise to future generations. Indeed, we cannot even imitate the conscious awareness of a dog or cat.

Why the human brain is so powerful

Our own brains contain around 86 billion neurons, 80 billion glial