

1

Introduction to AI

7

What is Artificial Intelligence (AI)?	8
What AI Can and Can't Do	9
Not One Size Fits All	10
Large Language Models (LLMs)	11
AI Jargon Explained	12
The Big Tech Battle for AI	14
The Future of AI	17

2

The Good, the Bad, and the Ugly of AI

19

The Human Factor	20
AI – a Force for Good	22
The End of Mankind?	23
AI Bias	24
AI Scams	26
Environmental Impact	30
Deepfakes	33
Less Human Interaction	34
Can you Trust the Technology?	35
Will the Robots Take Over?	36

3

Chatbots

37

What is a Chatbot?	38
Drawbacks of Chatbots	39
Accessing Chatbots	40
ChatGPT	42
Claude	54
Perplexity	58
Grok	62
More Chatbot Apps	66

4

System Chatbots

67

About System Chatbots	68
Copilot	70
Copilot App	78
Google Gemini	80
Apple Intelligence	84
Smart Speakers and AI	87

5

Photos and Videos

89

AI Image-Editing Apps	90
Removing Unwanted Objects	91
Night Mode	93
High Dynamic Range (HDR)	94
Special Effects	95
Images to Video	96
Firefly	98
Adobe Express	102
More AI Photo and Video Apps	106

6

AI in the Online World

107

AI on the Web	108
AI Search Agents	109
AI in Web Browsers	110
Dedicated AI Web Browsers	116
Email and AI	118
Social Media and AI	120
Social Media AI Images	121
More AI Online Apps	122

7

Writing

123

Too Good to be True?	124
AI Writing Options	125
Content Creation	126
ParagraphAI	130
Sudowrite	134
Translation	142
Copyright Issues	143
More AI Writing Apps	144

8

Family and in the Home

145

Getting a Taste for AI	146
Organizing with AI	148
Genealogy	151
Emergency Planning	152
Location Tracking	153
Heating	154
Lighting	155
Cleaning	156
Interior Design	157
More AI Family Apps	160

9

Entertainment

161

AI in the Creative World	162
Personal Recommendations	163
AI Characters	164
Making Music the AI Way	166
More AI Entertainment Apps	170

10**Health and Safety****171**

Detecting Health Issues	172
Apps for Symptom Checking	173
Health Monitoring Apps	174
Fall Detection	175
Sleeping on It	176
Smartwatches	177
Health and Fitness	178
Home Security	179
More AI Health Apps	180

11**Finance****181**

Behind the Scenes	182
Personalized Banking	183
Fraud Prevention	184
Stock Portfolio Predictions	185
More AI Finance Apps	186

Index**187**

1

Introduction to AI

Artificial Intelligence

*(AI) is now very much
part of everyday life.*

*This chapter explains
some aspects of AI and
looks to its future, too.*

- 8** What is Artificial Intelligence (AI)?
- 9** What AI Can and Can't Do
- 10** Not One Size Fits All
- 11** Large Language Models (LLMs)
- 12** AI Jargon Explained
- 14** The Big Tech Battle for AI
- 17** The Future of AI

What is Artificial Intelligence (AI)?

If one hundred members of the public were asked to describe Artificial Intelligence (AI) there would undoubtedly be a wide range of definitions and probably a few blank faces, too. Even computer scientists and researchers in AI sometimes struggle to come up with a unified definition. In general terms, it is the branch of computer science that designs machines (usually computers or robotic devices) that can perform tasks at a level matching, or potentially surpassing, human intelligence.

The birth of AI

Although the term “Artificial Intelligence” wasn’t coined until 1956 by a group of computer scientists led by John McCarthy at Dartmouth College in the US, one of the earliest examples of potential AI was in 1950. This was when the famous British scientist, mathematician, and code-breaker, Alan Turing, posed the question “Can computers think?”. He subsequently amended this to the idea of whether computers could imitate humans, leading to the Turing Test – see the tip.

During the second half of the 20th century, AI had a number of peaks and troughs, but it was still largely the domain of computer scientists, researchers, and AI enthusiasts.

AI comes into the limelight

At the beginning of the 21st century, some of the most talented computer scientists and researchers were working on AI, achieving some notable breakthroughs such as the stunning events of a computer beating a world champion in the game of Go, in 2016. But one of the most significant developments was when it was recognized that AI computers could be “trained” on large amounts of data, known as Large Language Models (LLMs). This led to dramatic progress that resulted in the release of the AI ChatGPT, which transformed the popular perception of AI.

In many ways, AI is still very much in the infancy of its development, with many opportunities and challenges ahead. Much of the technology has not even been invented yet, but the emergence of chatbots such as ChatGPT has given us a glimpse into the future, with their powerful application of AI.



The so-called Turing Test was posed by Alan Turing in the 1950s (he called it “The Imitation Game”, which also spawned the 2014 film of the same name). Essentially, the test is whether a human can tell the difference between communicating with a human or a computer. If they cannot tell the difference, then the conditions of the Turing Test are said to have been met. This test lives on, on the internet, with the use of the common CAPTCHA (Completely Automated Public Turing test to tell Computers and Humans Apart) test, used to identify if a user is a human or a robot.

What AI Can and Can't Do

Data is crucial to AI, whether this is text, images, videos or speech. AI specializes in analyzing and learning from data, to produce the output with which we are becoming increasingly familiar. But it is important to recognize AI's abilities and also its limitations.

What AI can do

As far as the output from AI is concerned, the way in which it can use data is widespread and dramatic. This includes:

- **Data analysis.** This is at the core of AI, and the results can be used in numerous different industries, from medicine to finance.
- **Creating content.** AI excels in this, whether it is replies to queries, writing a novel, creating artwork or making music. Chatbots are the most obvious example of this, but there are numerous dedicated AI apps for different types of content creation.
- **Automating tasks.** Repetitive tasks are ideally suited to AI, and it is an area where it has the potential to dramatically revolutionize the workplace.
- **Making recommendations.** Online and in the world of entertainment, AI can analyze user behavior and recommend specific content based on this.

What AI can't do

Despite some claims to the contrary in the past, AI cannot “think” like humans, nor does it have a conscience. It has no inherent sense of right or wrong, and does not have a moral compass for making decisions. AI should not be thought of as human, but capable of undertaking some human tasks. This is not to say that at some time in the future AI computers and robots will never have these attributes (and many computer scientists and researchers believe that this is a very real possibility; it is just a question of when this will be achieved). However, this is still some way in the future and, at present at least, there are still certain areas in which humans have an advantage over AI.



When thinking about AI, it is sensible to consider its output rather than the underlying technology, which is incredibly complex and sophisticated.



Once AI reaches a level of ability that matches humans – known as Artificial General Intelligence (AGI) – this is when it may be able to design future systems with far more advanced abilities than at present. Some of these systems may be able to display human-like qualities.



Although ChatGPT has been the dominant AI chatbot since it was released, the company behind it, OpenAI, released a “code red” warning in December 2025, identifying the fact that ChatGPT was falling behind some of its rivals. This highlights the highly competitive nature of the chatbot market and the importance that companies are placing on being ahead in the AI race.

Not One Size Fits All

When ChatGPT burst into the world in November 2022, it was many people’s first experience of consumer-level AI. Reassuringly, it also came in familiar formats, as an app and a website. Although apps are a relatively recent digital development, having been widely introduced with smartphones, they are now the main way of accessing digital content, from web browsers to social media sites and productivity options. So, when ChatGPT appeared as an app icon it was not an alien concept; the big change was the functionality and capability of this exciting new option.

Predictably, numerous other similar apps have followed ChatGPT. These are known as generative AI chatbots, which operate by being trained with vast amounts of text and images so that they can produce replies to almost any question and interact with human-like conversations. While ChatGPT has a significant proportion of the generative AI market, other popular chatbot apps include Claude, Google Gemini, Perplexity, Meta AI, and Grok.

AI everywhere

Despite the meteoric rise of AI chatbots, and their increasing popularity, it would be wrong to think that this is the only way in which we currently interact with AI. It is already embedded in a wide range of organizations and industries: banking, retail, medicine, public relations, law, environmental, transport, entertainment, sport, and government, to name a few, which are all heavily involved with AI. Most people will interact with AI on a daily basis, even if they don’t know it. AI technology similar to that used for chatbots is tailored for specific tasks and added to existing systems, or used to build entirely new ones. Some examples of this are: medical scans that are analyzed significantly quicker than by a human; banks using AI to identify the best financial products for their customers; media companies producing suggestions for people based on AI recommendations; and transport companies streamlining their systems for more efficient logistics. AI is spreading into all areas of society, and not just in app form.

Large Language Models (LLMs)

Large Language Models (LLMs) are the foundation on which AI chatbots such as ChatGPT are built. They consist of vast amounts of data (primarily text) on which the chatbots are then “trained” so that they can identify patterns and provide answers to a myriad of queries. Although they can provide very human-like responses, it is important to remember that this is done by analyzing the data in the LLMs: chatbots respond as computer programs, not humans.

Time for T

A key breakthrough in the development of LLMs was when researchers at Google found a way to effectively link pieces of text together, known as a transformer (the T in the acronym GPT, which stands for Generative Pre-trained Transformer). While Google was initially slow to develop this breakthrough, other companies were not so reticent, resulting initially in the launch of ChatGPT, with many other generative AI chatbots following.

It’s all about the data

As LLMs evolved, developers soon realized that the maxim “bigger is better” certainly applied in this case. The more content that was fed into an LLM, the more efficient it became. This raised issues about sourcing the content and the amount of energy and computing power that was required to create and maintain LLMs but, in general, the AI industry has brushed these aside in the quest for ever-larger models to power their chatbots.

Legal woes

The development of LLMs has not been without controversy. The technology is ravenous for data, and this has led to companies using much of the content on the internet (particularly text, images, and music) to populate LLMs. Taking data from the internet in this way is known as “scraping”, and it is frequently done without the permission of the websites or content owners from where the data is taken. This has led to numerous court cases for copyright infringement, with many of them still ongoing.



Chatbots are known as generative AI chatbots, through their ability to generate content using AI.



Although Google was not the first company to release a commercially successful chatbot, they have since caught up in spectacular fashion, and their Gemini 3.0 version is now one of the most popular chatbots on the market.

AI Jargon Explained

As with any new and developing technology, AI is awash with jargon and acronyms. Some to look for include:

- **AI algorithms.** These are sets of computing rules that are used to enable AI agents (either machines or apps) to use data to replicate human intelligence.
- **Artificial General Intelligence (AGI).** This is thought to be the next step for AI, and it is when AI agents are considered to be able to do all of the tasks of humans, rather than just individual tasks. AGI is also known as human-level intelligence.
- **Artificial Superintelligence (ASI).** This is the level beyond AGI, at which point AI agents will be significantly superior in their abilities to humans. If this is achieved, it is unknown how quickly multiples of this superiority will be achieved.
- **Chatbots.** These are AI agents, usually in the form of apps, that can be used to conduct human-level conversations, either by text or speech, to perform numerous tasks, including answering queries, manipulating and creating images, creating text, and writing computer code.
- **ChatGPT.** The most widely used chatbot, created by OpenAI and one of its founders, Sam Altman. At the time of printing, there are approximately 800 million active users interacting with ChatGPT each week.
- **Data centers.** These are large collections of computers that are used to handle AI requests. There is massive investment in data centers by AI companies, as they are crucial to the development of the technology. Some data centers take up the space of several football fields, and they use vast amounts of energy.
- **Datasets.** These are structured collections of raw data that can be used to train LLMs (see the next page) that are used by AI chatbots.



Artificial Superintelligence (ASI) is also sometimes known as just Superintelligence.

...cont'd

- **Generative AI.** This is AI technology that uses Machine Learning (ML) to create content used by chatbots, such as text, images, and videos.
- **Generative Pre-trained Transformer (GPT).** This is the process by which chatbots interpret the data in LLMs to produce results for users.
- **Hallucinations.** This is a very human term that refers to the fact that generative AI chatbots can sometimes produce answers that are inaccurate or nonsensical.
- **Large Language Models (LLMs).** These are huge collections of data that are then used to train chatbots.
- **Machine Learning (ML).** This refers to computers using large amounts of data to “learn” certain tasks, such as playing chess. The AI function predicts certain outcomes based on the data, rather than being programmed for every task.
- **Natural Language Processing (NLP).** This is the technology that helps to power chatbots, by processing text and speech to generate human-like communication.
- **Neural Networks (NNs).** This is an area of AI that supports Machine Learning and is based on the structure of the human brain, but is not intended to replicate it.
- **Organoid Intelligence (OI).** This is a developing technology that aims to combine human biology with computer systems.
- **Paperclip Maximizer.** This refers to a thought experiment conducted by the philosopher Nick Bostrom that was designed to show potential unintended consequences of AI agents – see the tip.
- **Unintentional bias.** This is the tendency of AI to replicate racial and gender biases in society. It is a significant issue for AI.



The paperclip thought experiment, known officially as the Paperclip Maximizer, was devised by the philosopher Nick Bostrom in 2003. The premise is that an AI agent is given a single task of producing paperclips. Initially it does this with the correct raw materials, but when these run out it turns to other materials, since its only focus is to fulfill this task. In this way it destroys the planet, and everyone in it. The experiment highlights some of the potential consequences of AI if it is not designed, constrained, and regulated properly.



Since April 2023, DeepMind has been known as Google DeepMind.



Much of DeepMind's development of AI was based on Reinforcement Learning (RL), training AI agents to play games based on data from previous experiences. One of the greatest successes in relation to this was the AI program AlphaGo, which was trained to play the notoriously fiendish game Go. In 2016, AlphaGo beat a Go world champion, Lee Sedol, with one move in particular showcasing the remarkable power of AI, and stunning the watching audience.

The Big Tech Battle for AI

The history of AI goes back to the 1950s, but this century two companies have shaped the AI landscape. OpenAI was founded in the US in 2015 by, among other people, Sam Altman, and DeepMind operated in the UK, from 2010, with its founders Demis Hassabis, Shane Legg, and Mustafa Suleyman. Both companies harbored ambitions of developing Artificial General Intelligence (AGI), which is a level of AI that operates at an equal or higher level than humans for all of the tasks that humans can perform. Although the approaches of OpenAI and DeepMind were different, they both made significant breakthroughs within a few years of being founded.

Both companies quickly attracted the attention of big tech, keen to use AI to bolster their profits. Initially, the founders of DeepMind and OpenAI were reluctant to align themselves with any of the big tech companies, as they feared that their altruistic aspirations of helping mankind would be diminished in favor of improving the balance sheets. But the inclusion of big tech was inevitable due to two things: money and power (computing power). The wage bills at OpenAI and DeepMind were astronomical and, in order to keep the best and brightest in the business, the companies had to pay top dollar or risk losing their workers to competitors, such as the big tech companies themselves. Luckily, most of the staff believed in the AI projects that they were working on, but they still had to be paid well. Another cost was in the technology of developing AI models. These require huge amounts of processing power and cloud storage, all of which comes at a high price. The founders of DeepMind and OpenAI also realized that their AI models got better as they got bigger, requiring more computing power and storage. To progress, massive investment was needed, and it had to come from somewhere.

Compared to the big tech companies, DeepMind and OpenAI were in no position to expand their work to the levels that were required, or, if they did, it could take decades to try to achieve their goals.

...cont'd

Acquisition of DeepMind

Even though one of the early investors in DeepMind was Elon Musk, by 2013 the company realized that they needed a more substantial investment to continue on the path that they had outlined. A potential investment from Mark Zuckerberg's Facebook fell through in 2013, reportedly due to the insistence from DeepMind's founders to include an ethics board in the deal to oversee the future direction of AI development. A year later, at the beginning of 2014, a deal was agreed with Google to acquire DeepMind for approximately half a billion dollars. Although the work of the company initially continued in the same general direction, along with their much desired ethics board, it was very much now in the hands of big tech.

Microsoft buys into OpenAI

OpenAI followed a similar but not identical path to DeepMind, seeking investments rather than selling the entire company. This attracted the attention of the tech giant Microsoft, which invested \$1 billion in 2019. This led to a very close working relationship between the two companies, and Microsoft invested a further \$10 billion in 2023.

ChatGPT changes the script

Helped by the investment from Microsoft, OpenAI released the ChatGPT chatbot in 2022. Almost instantly, ChatGPT became the identifiable manifestation of AI for the consumer market. This was no longer science fiction perused by computer geeks; this was a digital revolution where you could interact with an app as if it were a human and generate remarkable results for almost any query imaginable. ChatGPT stunned the computing community and enthralled the general public. Some of the people who were most stunned by ChatGPT were the executives at Google. For at least two decades they had enjoyed almost total dominance in the area of web search and the vastly lucrative business of online advertising that accompanied this. The ease by which people could obtain results from ChatGPT seemed like it could pose an existential threat to Google.



The corporate structure of OpenAI has become increasingly complex over the years.



Parmy Olson's excellent book *Supremacy* provides a compelling insight into the battle to gain an advantage in the AI race, primarily between Google and Microsoft.



Although public attention concerning AI is largely taken up with chatbots such as ChatGPT and Gemini, much of the most profitable areas are in the construction of data centers and cloud storage systems. Generating profits from chatbots is still something of a challenge for the AI industry.

...cont'd

Google embraces generative AI

The response by Google to the threat of ChatGPT was to speed up development of their own generative AI chatbot. Ironically, part of the technology used in chatbots (the transformer element of GPT) was first developed by researchers at Google, but the company was reluctant at the time to deploy anything that might negatively impact their online ad (advertisement) business. The result, after a few incarnations, was the Google Gemini chatbot. In November 2025, Google released Gemini version 3.0 as a direct competitor to ChatGPT, as the race for AI domination continues.

Big tech is all around

Although Microsoft and Google were two of the companies that invested at an early stage in technologies now widely used in AI chatbots, this does not mean that other technology companies have been standing idly by. Nvidia Corporation has become the world's largest company in terms of market capitalization by powering AI revolutions with, among other products, the processing chips that are used to run increasingly power-hungry AI computers. Meta (the parent company of Facebook, Instagram, and WhatsApp), Amazon, and Apple have all also invested very heavily in areas of AI, from data centers to cloud storage. This has helped to centralize the control of AI development to a small number of very rich companies. How this will influence the development of AI is difficult to predict, but these companies will certainly expect an impressive return on their considerable investments.

It is not just in the US where big tech is investing vast sums of money in AI. In China, there are several companies that are leading the way in the AI space, and they could become serious competitors to their American counterparts. Alibaba, Baidu, and Tencent are emerging as significant players in the AI marketplace, and it is likely that some impressive innovations will emerge from here, as it did with the release of the DeepSeek chatbot in 2023 – see the next page.

The Future of AI

AI has advanced at one of the fastest rates of any technology in the history of mankind. It also has the potential to have the greatest impact, beyond that of the Iron Age, the Industrial Revolution, and the internet.

US and China

When considering the future of AI it is important to be aware of the geographical and geopolitical issues. AI development is most advanced in the US, but China is becoming a major player in this area and is investing heavily in all aspects of AI technology. This has led to something of an AI arms race between the world's two largest economies. The dominance of the US in the future of AI is not guaranteed, which was notably demonstrated in 2023 when the Chinese chatbot DeepSeek briefly took the AI world by storm. Since the US and China have differing views of the world order, how they both develop AI in the future will have an important influence on many areas of global society. In addition, many other nations, such as those in Europe, India, Japan, and Korea, will be looking to develop their own AI industries, with the consequences not always being clear.

Money

AI is already a trillion-dollar industry, when you consider the amount of investment in processing chips, data centers, cloud storage, energy generation, and personnel. Fears have been raised that this is creating a financial bubble, in the same way as the early development of the internet created the dotcom bubble, which burst spectacularly at the beginning of the 2000s. One area of concern is known as circular funding, whereby different elements of the AI industry invest in each other, without any of the money dispersing to the wider economy. Given the amount of money being spent on AI, the question of a bubble is justified, as are the voices that raise concerns about how AI is going to generate these levels of revenue in the future. Advertising is one of the most likely areas of revenue, with highly targeted AI adverts, based on vast amounts of personal data that can be harvested by AI.



An example of circular funding is where an AI software company may invest in a chip processing company, which in turn invests in data center companies, which then invest in software companies.



When AI is considered to be performing at the level of humans, this is known as Artificial General Intelligence (AGI). The level above this is known as Superintelligence, coined by the philosopher Nick Bostrom, when AI can dramatically outperform humans. It is at this point where AI may be able to develop itself without any human involvement.

...cont'd

AI developing itself

In terms of the future of AI, it is important to remember that it is still very much in its infancy. The potential for development is almost limitless, particularly when all breakthroughs and technologies that have not even yet been invented are taken into account. There is also a likelihood that at some point AI will surpass the level of human intelligence, meaning that it could have the potential to develop itself without the need for any human input. This is unknown territory and means that the AI world in decades to come will be dramatically different from the present day, with some results being hard to predict.

Containment

The fear of many people in the world of AI is that the development of the technology is proceeding at a breakneck speed, without enough effort being made to provide containment features to prevent it from running out of control. This can be done in the design of AI technology itself, although there can be a reluctance to do this robustly enough when there are huge profits at stake. Also, there is a role for national authorities and governments to play, to put in place legislation and policies that restrict any potential rogue uses of AI. Unfortunately, legislation of this nature is hard to agree upon even on a national level; far less on a global one. The development and expansion of AI currently far outweigh any efforts to contain it.

The AI future is ours

The future evolution and development of AI is both extremely exciting and deeply worrying. The challenge, on an individual, national, and international level, is to navigate a path that ensures as many of the benefits of AI for mankind as possible, without falling foul of too many of the pitfalls. It is not an easy task, and it will take clear-minded thinking, innovation, compromise, cooperation, and not a little luck.

The future of AI, and its impact on society and the world, is in our hands however we choose to tackle it.