INTRODUCTION

Introducing generative AI
and ChatGPT.

The dawn of the AI age

A computer did not write these words. But it easily could have. As you probably know, we now have technological tools that can generate text instantly, on request, that seems equal to anything a human could produce – at first glance, at least.

The term for this type of technology is *generative artificial intelligence*, or generative AI, because it can 'generate' content like text, artwork and music in response to a user's text prompt. Popular generative AI writing tools include OpenAI's ChatGPT, Anthropic's Claude, Google's Gemini, Meta's Llama and Microsoft's Copilot. In this book, I will refer to ChatGPT alone because it's the most popular and well-known, and the others offer broadly similar functions.

Generative AI feels like a seismic shift in our relationship to technology. For the first time, computers are laying claim to cultural production, and even artistic expression, that has previously been done by humans alone. For decades, we've had computers that can do boring science stuff like calculating, planning and data-processing. More recently, we've had computers that can search the whole of human knowledge at the touch of a screen, keep us in constant contact with everyone

we know and cater to our tastes before we even consciously decide what we like. But we've never had a computer that can come up with something like this:

> From circuits born, a mind begins to dream,
> New worlds arise from code's creative stream.
> Art, thought, and form, reshaped by unseen hands,
> A future forged where human and machine stands.[1]

The reaction to the appearance of generative AI has been frenzied. Businesses, investors and the general public have all piled in, lured by the promise of the 'next big thing'. Some creatives have been understandably fearful that they are about to be replaced by a machine – yet others have willingly embraced the technology and encouraged others to do the same. And while some observers have hailed generative AI as the dawn of a new age, others have been more critical of the technology and the buzz that surrounds it.

With so many people shouting so loudly about AI, it can be hard to tell fact from hype. In this book, I've aimed to take a realistic look at what writing tools like ChatGPT can and cannot do, particularly in relation to the craft of writing as practised by human writers. But first, let's find out where these AI writing tools came from.

A brief history of ChatGPT

OpenAI was founded in 2015 by Sam Altman, Elon Musk, Ilya Sutskever and Greg Brockman. In its early days, it was a non-profit with a mission to 'advance digital intelligence in the way that is most likely to benefit humanity as a whole'. By 2018,

the firm was experimenting with Generated Pre-Training, or GPT. In the following year, OpenAI struck a deal with Microsoft to accept $1 billion of funding and use Microsoft's Azure cloud computing platform, at which point OpenAI switched from being a pure non-profit to a complex web of interlinked companies with Microsoft as minority shareholder.[2]

Early GPT models didn't cause much of a stir outside the tech community. But all that changed when ChatGPT appeared in November 2022, using the GPT-3 model. Taking the form of a chatbot, which interacts with users through typed conversations, ChatGPT offered an impressive range of functions. It could research topics and summarise its findings. It could recommend books, films, hotels, music and more. It could perform calculations and write computer code. It could edit and proofread existing text. And, most importantly for our purposes, it could generate plausible prose, poems, stories, speeches, marketing copy and news articles – all in a style of the user's choice.

News of the remarkable new generative AI tool spread like wildfire, and just two months later, ChatGPT reached over 100 million users, making it the fastest-growing consumer software application of all time. In March 2023, OpenAI launched GPT-4, which offered many improvements over GPT-3 and was available through a paid subscription. As I write this, ChatGPT is among the top 20 most visited websites in the world.

Investors, sensing a historic opportunity, have been pouring cash into AI. In 2023, the industry was already worth around $200 billion, and some expect it to mushroom to over $1.8 trillion by 2030.[3] Meanwhile, businesses have been eagerly exploring the possibilities of using AI to automate their

processes, discover valuable new insights or build their own chirpy chatbots to augment customer service teams.[4]

Given AI's stratospheric rise to date, it's hardly surprising that there have been some eye-popping predictions for its future. Apparently, by 2030, we'll all be using AIs as our personal assistants, tutors, careers advisors, therapists, accountants, lawyers and even romantic partners,[5] while physical robots performing human tasks will be commonplace in our homes and neighbourhoods.[6] Sam Altman, CEO of OpenAI, predicted that AI would probably replace most of the work done by humans, while also heralding a massive increase in 'material abundance' through 50 percent year-on-year increases in global GDP.[7] If we're lucky, AI will even step in to fix climate change[8] – which is only fair, since AI blazes through up to 33 times as much energy as regular software to complete the same task.[9]

Such brash optimism is typical of Silicon Valley, with its philosophy of 'move fast and break things' and a business model that depends on projecting a sense of rapid progress and urging people to get involved by cultivating FOMO – the fear of missing out. However, not everyone is feeling quite so positive about the prospects for AI. Even Altman himself has confessed to feeling 'a little bit scared' about its future, admitting that AI technology could be used for cyberwarfare, hacking, disinformation or biological warfare – particularly once authoritarian regimes get up to speed with it. He's also speculated that AI could end up making us stupid, by robbing us of our 'discriminating thought process'.[10]

By 2024, a 'vibe shift' seemed to be building over AI. OpenAI continued to release new versions of ChatGPT, including one that could speak as well as write, but still could

not explain precisely how it planned to recoup the exorbitant costs of building and training AI models.[11] Some sources claimed that ChatGPT's growth was flatlining as users lost interest.[12] Sceptics pointed out that AI applications vastly outnumbered actual use cases, stoking suspicions that AI was a solution in search of a problem.[13] Investors, no longer quite so exuberant, grew impatient for results, with many fearful that AI would turn out to be the biggest tech bubble of all time.[14]

As I write, the story is still unfolding – but I think we can confidently say that the first era of uncritical enthusiasm for AI is over. As with previous 'next big things' like wearable technology, NFTs, bitcoin and the metaverse, AI may be settling into place as 'just another technology': important and useful in its way, but not quite as earth-shattering as we first thought. But even after the hype does subside, the tools themselves will still be around, and we'll still have to decide whether and how to use them. So what can ChatGPT actually do? And how exactly does it do it?

How ChatGPT works

ChatGPT is a Large Language Model (LLM). It works by recognising and reproducing patterns in human language through what is called Natural Language Processing, or NLP.

Before ChatGPT was launched, the GPT-3 engine had to be 'taught to speak'. To achieve that, it was 'pre-trained' on 570 gigabytes of text available on the internet in the form of websites, books, news articles, journals and social media. (By my calculation, that's around 94.6 trillion words.)

By analysing these myriad sentences and exploiting the statistical regularities in language, GPT-3 was able to learn the 'rules' that govern how human beings speak and write. As it wandered through the universe of language, it observed which words tend to be used together, which ones are more important and which ones are most likely to come next in a given sequence.[15] Essentially, GPT-3 is a word-prediction machine.

> You give [it] a text and cover up one of the words. Then you make it guess the missing word. Then you do it again. And again. And again. Over time, feeding it with new texts and more data, it does not only get increasingly more accurate at guessing words, but it also gains a vast amount of knowledge and learns to produce humanlike texts and responses.
>
> DANIEL HARDT[16]

Natural Language Processing is what allows ChatGPT to hold humanlike conversations with its users. You can type in a request (known as a 'prompt') that gives instructions to the software in standard English, as if you were talking to another person, and ChatGPT will respond in kind.

For example, you can:

- Ask ChatGPT open questions, like 'Tell me about the history of breadmaking'
- Ask ChatGPT to write text, like 'Write a sonnet about the joy of breadmaking' or 'Write a job application to become a music teacher'

- Ask ChatGPT to edit text you already have – for example, 'Rewrite the text below in plain English' (and you paste in the text)
- Ask ChatGPT to summarise your text, or even someone else's – for example, 'Write a plot summary for *Wuthering Heights*'
- Hold a conversation with ChatGPT – for example, by asking it a follow-up question or telling it to alter or refine the text it has just produced. ChatGPT automatically 'remembers' what you're talking about from one prompt to the next.

Arthur C. Clarke said, 'Any sufficiently advanced technology is indistinguishable from magic.' That aptly describes the experience of seeing ChatGPT in operation for the first time. As the screen rapidly fills up with logical, grammatical and plausible content at superhuman speed, it seems that this machine has mastered a skill that takes us years to learn and remains a challenge for the rest of our lives: writing. And you can chat to it too.

Take a closer look, however, and it's clear that ChatGPT has its problems. In terms of factual writing, it makes errors (known as 'hallucinations') that may be difficult to detect. Its writing tone can sometimes sound weird or robotic, and it lacks empathy or insight into human life. Morally speaking, the whole enterprise is questionable, as ChatGPT's capabilities are built on the past labours of human writers who receive no payment or even credit for their work. And worst of all, the platform has negative impacts on the user, by taking away a task that they can and (I believe) should do for themselves.

Just a Luddite?

I'll be expanding on all my criticisms of ChatGPT in the second part of the book. But before we go any further, I have to pause and answer the charge that I am nothing more than a Luddite – that is, someone who resists technological innovation not on its merits, but in a stubborn or reactionary way.

Behind this accusation lies the assumption that technological advances always mean societal improvement. New technologies expand the range of possibilities by allowing us to do more things, or better things, or the same things more quickly and easily. As our knowledge expands outward, so our standard of living improves. Essentially, new tech means a better life.

Obviously, some new technology *does* bring benefits. Advances in medicine save lives. Improvements in agriculture prevent famine. Better transport saves time and fuel. And digital technology can manipulate knowledge and information in remarkable ways. But those benefits might still come at a cost, and that cost might not become clear until the technology has been widely adopted and used for some time. For example, the internet has brought us new connections, saved us time and made human knowledge freely available to all. But it's also sown discord, spread misinformation and made our children more anxious than they've ever been.[17]

So when a new technology like AI comes along, we're entitled to ask, in a polite and measured way, what the downsides might be – particularly when we're bombarded with so much breathless coverage of the upsides.

Some AI evangelists argue that anyone who expresses reluctance is suffering from fear of the unknown, fear of change or fear of becoming irrelevant. This frames the discussion in

terms of feelings vs facts: our irrational fears are holding us back, and once we overcome them with a dose of factual knowledge, we'll be able to move forward.[18] A cruder accusation is that Luddites simply don't understand the technology they're resisting. If only we'd make the effort to learn a bit more about it, we'd see how beneficial it can be. Or maybe we're simply incapable of learning, in which case we probably deserve to be left behind.

When you face a big decision with major consequences, it's perfectly sensible to take a step back, wait for more information and weigh up benefits versus costs. That's not fear, but rational caution. Then, having learned more about the possibilities, you might still make a rational or principled decision not to get involved. And on top of all that, if you feel an instinctive unease that something is not quite right, or seems too good to be true, it often pays to act on it, or at least look into it.

The original Luddites smashed new automated machinery introduced into English textile mills during the Industrial Revolution. Over time, their name has become a byword for anyone who resists new technology or even progress of any sort. But the Luddites were neither fearful nor ignorant of industrialisation, and nor did they have a knee-jerk opposition to anything new. Rather, they took a clear-eyed view of the social impact of technology and correctly predicted what would happen when industrial automation met the power of big capital in 19th-century English society. While smaller textile mills might well be run with respect for workers' wellbeing and the law of the land, larger enterprises would have enough clout to disregard the rules. That would mean lower pay and worse conditions for workers, reduced quality for customers and colossal profits for those who owned the mills.

> The Luddites knew exactly who owned the machinery they destroyed. They saw that automation is not a faceless phenomenon that we must submit to. And they were right: Automation is, quite often and quite simply, a matter of the executive classes locating new ways to enrich themselves.
>
> Brian Merchant[19]

It wasn't the machines themselves that Luddites had a problem with, but the power of the companies that built them. Smashing the machines was merely the only means at hand to confront and challenge that power. And when factory owners mobilised military and legal force to crush the Luddites' protests, they were proved absolutely right.

In the same way, questioning the value of generative AI is not about being fearful or dismissive of the technology itself, or new technology in general. It's about resisting the powerful corporations who offer to solve non-existent problems, or proclaim that everything can be quicker, easier and smoother, but without ever explaining *why*.

Tech critic Evgeny Morozov calls this 'technological solutionism': the notion that technology can solve all the problems of humankind, and that once engineers have built the right solutions, we'll enter a frictionless, hassle-free utopia.[20] Each new digital platform arrives with promises of an easier and better life. But after the initial thrill of novelty and the rush of early adoption, we soon end up mired in what Cory Doctorow calls 'enshittification': the slow degradation of an online product or service as it surreptitiously shifts from attracting users to monetising them, often through

advertising.[21] Ultimately, tech firms undermine the core user experience by prioritising shareholder value above all else.

ChatGPT is only at the beginning of that process, but the direction of travel is already clear – and the benefits are dubious even now. And that brings us to the second part of the book, where I lay out my case against ChatGPT. Then, in the third part, I'll suggest some ways to make your writing more human, stand out from machine-generated content and resist the advance of AI.

THE CASE AGAINST CHATGPT

Why a machine will never
write as well as you.

The process of writing

To set the scene, let's define some terms. When I say the word 'writing', you might think of using a pen and paper or typing at a keyboard. But there's a lot more to the writing process than putting words onto a page or screen.

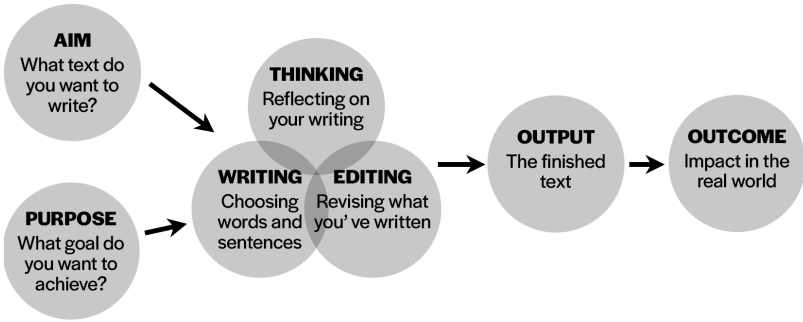

First, you decide what you're going to write – your *aim*. For example, you might be planning to write a crime novel. Or you might be writing an article about solar panels aimed at homeowners who are considering installing them. Or you might be working on a vital email to your boss about an upcoming job opportunity.

At the same time, you decide the *purpose* of your writing. While your aim captures what you want your writing to *be*, your purpose describes what you want it to *do*.

Crucially, your purpose should be framed in terms of your reader(s), not you as a writer. It will usually be making your reader know, think or feel something new or different. For example, with a crime novel, your purpose is to entertain, intrigue and thrill your reader with the twists and turns in your story. For the solar panels article, you want to provide clear information or valuable guidance to support a decision to buy. For the email, you want your boss to understand why you'd be perfect for the role and offer you an interview.

Once you've decided on your aim and purpose, you move on to *thinking* about what to write. This could include writing a plan, like the plot of a novel or the outline of a non-fiction book. You might also do some research – for example, you might delve into the historical setting for your novel, investigate the legal regulations on solar panels or ask around the office to find out what your company looks for in a manager. Or you might just reflect on your writing project in a less formal, structured way, thinking it over in your mind and jotting down ideas or phrases.

Next is the actual *writing*: the act of choosing words and putting them down. As you know, this can prove extremely challenging, as you struggle to find the words that will achieve your purpose and realise your aim. Even if some parts of your writing take shape relatively easily, you'll probably run into problems at some point. And of course, it's precisely this frustration that ChatGPT offers to take away.

After writing comes *editing*, when you're mainly trying to improve what you've already written rather than produce

anything new. This can involve changing individual words, reordering sentences or paragraphs, deleting bits that aren't working and so on.

Thinking, writing and editing all overlap. You'll often move back and forth between them as your writing develops, and it can be hard to say where one stage ends and another begins. For example, few writers generate an entire first draft before editing a word; it's more common to make changes as you go along, switching between editing and writing from moment to moment. And thinking is always there in the background: you'll probably carry on mulling over your writing, whether consciously or unconsciously, even when you're not physically at your desk – the mental process known as *incubation*.

The thinking-writing-editing phase is fairly open-ended, and continues until you either reach a version you're happy with, run out of time or simply give up, at which point your writing is done. The end result is the *output* of your writing process, and it reflects what you targeted when you chose your aim (a book, an article, an email, etc).

However, that's not the end of the story, because your writing must now go out into the world, be read by its readers and achieve the purpose you set at the beginning. This is the *outcome* of your writing.

Outcomes are the true test of writing success. Unless you're writing something purely private and personal, like a diary, you want your writing to get some sort of real-world result, and it only succeeds insofar as it does so. If your writing doesn't do what you wanted it to do, it's failed – and all your efforts have gone to waste.

Overall, writing is a *process of communication* that links the writer and the reader. Your goal is not just to 'write something'

or 'create some content', but to communicate ideas, facts, stories, opinions or emotions and make something happen in the mind and heart of your reader. The words on the page are just the tool you use to make that happen.

Now compare this process to what happens when you use ChatGPT to produce your text instead:

- While ChatGPT offers instant results, the prompt/response format neglects vital aspects of the writing process. Your prompt is focused mainly on the output you want to obtain, as opposed to the purpose you want to achieve.
- Although you can try to explain your purpose to ChatGPT, it has no objective knowledge of the world and doesn't understand human experience, goals or emotions. So while it's relatively easy to change the structure or style of the output, ChatGPT will never comprehend the *substance* of what you're trying to say or what it will mean to your reader.
- ChatGPT offers to do the thinking for you, but you have no idea what it's doing, or how. All you can do is look at the content it creates and infer its 'thinking' process from that. You also miss out on doing the thinking yourself, which can have a major impact on both the direction and quality of your writing.
- ChatGPT seems to offer 'the answer' to your prompt, but its output is only one possible approach among many.
- ChatGPT seems authoritative, but it can make mistakes, and it can be hard to detect minor errors if you don't know the subject well.

- Last but by no means least, ChatGPT cannot write from your mind and heart, drawing on your own absolutely unique experience and view of the world, because that is something that only you can do.

In the rest of this chapter, I'll explore these areas in more detail, getting deep into the whys and wherefores and citing evidence to support my points.

ChatGPT is not a person

For as long as computers have existed, both their creators and their users have tended to anthropomorphise them – that is, to see them as somehow human. The ground was prepared by popular science fiction, where we've met the malignant Hal 9000 in *2001: A Space Odyssey* (1968), the child-savant Brain in Isaac Asimov's robot stories[22] and the oracular Deep Thought in Douglas Adams' *Hitch-Hiker's Guide to the Galaxy*,[23] among many others. All these imaginary computers had distinctive characters, held humanlike conversations with the people who used them and were created to solve problems that humans could not.

In the real world, engineers do the same thing, giving their creations names like Siri and Alexa so we're encouraged to think of them as people. ChatGPT is an exception, although the word 'chat' still encourages us to think of our interactions with it as a human conversation.

With all these devices, we intuitively use speech as a measure of intelligence and an indication of humanity. When we come across something that speaks in the same way that we do, we assume it must think like us too. If Hal had communicated his