

# The Intelligent Interface




Scarcity in **Abundance**

# Table of Contents



04



10



14

- 04** – Introduction: The Intelligent Interface
- 10** – Chapter 1: Interface as Capability, Activity, and Manifestation
- 14** – Chapter 2: Errors in Judgement Are Timeless
- 23** – Chapter 3: The Prompt Factor



23



44



32

- 32** – Chapter 4: The Capability Factor
- 44** – Chapter 5: The Form Factor: The Next Big Thing
- 53** – Chapter 6: Conclusion and Speculative Outlook

# Introduction: The Intelligent Interface



The rise of artificial intelligence (AI) as 'the' new interface can be felt in many organisations. The use of AI for information gathering, also known as 'prompting' is becoming more commonly used. Prior knowledge is less important on the work floor as AI tirelessly provides relevant answers. In your organisation, you might see a discord between people who won't have any of it yet, on one side, and those who fully embrace AI, on the other. With the latter not fully committing as necessary guidelines, or 'guardrails', are still in development.

While goalposts are being set in boardrooms and the first experiments can be carried out, AI developments seem to move faster than ever. Organisations who are yet to start implementing AI strategies are at risk of playing catch-up. Amy Webb, 'Quantitative Futurist', makes this clear in her 'tirade', in which she speaks of an accelerated technological era, or a 'supercycle'.

According to Webb, we are at the dawn of a period of unprecedented development of AI technologies, giving a new dimension to the digital transformation of each organisation.

The concept 'supercycle' is based on the theory of economist Nicolai Kondratiev, who about a hundred years ago suggested that society develops in long waves of roughly sixty years. Important technological breakthroughs, such as the steam engine, mark the onset of such a wave. Webb challenges this theory, as technological developments unfold much more rapidly now than he proposed at the time. Webb, acknowledged as one of the most important tech influencers, frequently shares her vision in trend reports that are eagerly anticipated. In March 2024, she spoke at the prestigious South by

Southwest (SXSW) festival in Austin, Texas, where many gathered to learn about her vision for the future.

Amy Webb's story aligns with a wider trend in which technological progress and optimism play a central role. Recently, the term 'e/acc', which stands for effective accelerationism, has been mentioned more and more online. It stands for a movement that embraces the benefits of technological development, with Marc Andreessen, founder of Netscape and well-known venture capitalist in the tech industry, as one of the pioneers. In his profile on platform X, he uses the indication "e/acc" to emphasise his affinity.


For most, generative AI is mainly an online space where they can go to use prompts to ask questions that may or may not be answered adequately. But the assumption is that it will soon penetrate every fibre of the information society. We'll be able to see AI as the new operating system of organisations. With the launch of new devices that are based on Large Language Models (so-called connectables) and a new generation of robots animated with generative AI, this vision of the future is approaching at extraordinary speed.

**Webb and the 'e/acc' movement share a common view: the anticipation of a significant technological, economic, and social acceleration driven by recent breakthroughs in artificial intelligence (AI).**

AI is seen as a system technology (General Purpose Technology, GPT<sup>1</sup>), comparable to past revolutionary technologies such as railways or electricity. These GPTs, also referred to as meta technologies, are crucially important for broad-based economic growth. Artificial intelligence as the newest GPT promises to encourage a similar expansion and will become an integral part of our economic infrastructure.

---

<sup>1</sup> The same acronym as for Generative Pre-trained Transformer in ChatGPT, but not related.



Although these new devices admittedly use LLMs, the focus lies on their characteristic of being able to perform multiple tasks simultaneously and independently, as ‘software agents’ are working in the background. Those who have worked in IT for longer, might still know the word ‘webcrawler’: smart packages we sent online to arrange things for us. A dream that never fully became a reality but now has been given a second life. In short, artificial intelligence nestles as an intelligent interface between our physical and digital worlds, bringing an incredible variety of new possibilities. A new word has been invented for it already: ‘all-in-one-interface’. Bill Gates, founder and former CEO of Microsoft is excited about its possibilities. He writes: “You won’t have to use different apps for different tasks. You’ll simply tell your device, in everyday language, what you want to do.” In the near future, anyone who’s online will be able to have a personal assistant powered by artificial intelligence that’s far beyond today’s technology.<sup>2</sup>

---

<sup>2</sup> Gates, B. (2023, 9 November). AI is about to completely change how you use computers. *GatesNotes*. <https://www.gatesnotes.com/AI-agents>

This report explores developments that lead to that intelligent interface. It provides a status update on the most recent events in the area. We will place them in the context of the relatively young history of information technology. After all, according to experts, we have been in the information technology supercycle since the 1990s, delivering unprecedented acceleration of various processes. Adding an extra layer of AI on top of that supercycle should be seen as ‘the smartening of the acceleration’. Amy Webb’s straightforward conclusion is that we’re making everything that is dumb smart now. And Marc Andreessen, more cryptically, says:

“Let’s make sand think!”

Before you continue reading, we would like to draw your attention to two important conclusions from our earlier research. We refer to the two reports that we recently published under the banner ‘Scarcity in Abundance’. The first conclusion was that the impact of a new technology has several order effects. The order effect that Amy Webb and the ‘e/acc’ movement refer to is that of the acceleration of economic developments. But there are many more effects that may or may not have been anticipated. For a more detailed description we’d like to refer you to our report *The Great AI Experiment*.<sup>3</sup>

The second conclusion was that the physical and digital worlds are more and more at odds with each other. The depletion of the earth, loss of biodiversity, and global warming present us with tremendous and unprecedented challenges. Sam Altman, Ope-

---

<sup>3</sup>See <https://ict-books.com/book/the-great-ai-experiment/>.





nAI's chief, already warns us about a new energy crisis and ties the future of AI to breakthroughs in the field of clean energy delivery. At the World Economic Forum in Davos, he said that AI uses many times more energy than one thinks, and that nuclear fusion is a requirement to meet the enormous demands for energy.<sup>4</sup>


It would be good to realise that the technological optimism of Amy Webb and 'e/acc' are part of the firm belief that technology will save us from global ecological disaster. Marc Andreessen recently published his Techno-Optimist Manifesto, in which he amongst other things declares that the earth is dramatically *under* populated. The world's population could easily expand to fifty billion people, and we could ultimately explore other planets. In our report *The Rise of the Experience Ecology*<sup>5</sup>, we elaborate on this techno-optimist philosophy and present the alternative of a Post-Growth Society. A well-loved statement fitting with post-growth ideology is that of economist Kenneth Boulding:

**“Anyone who believes exponential growth can go on forever in a finite world, is either a madman or an economist.”**

---

<sup>4</sup>Reuters (2024, 16 January). OpenAI CEO Altman says at Davos future AI depends on energy breakthrough. <https://www.reuters.com/technology/openai-ceo-altman-says-davos-future-ai-depends-energy-breakthrough-2024-01-16/>

<sup>5</sup>See: <https://ict-books.com/book/the-rise-of-the-experience-ecology/>.




# Chapter 1: Interface as Capability, Activity, and Manifestation

An intelligent intermediate layer in our dealing with computers could be very useful, but please explain: what does that mean for me? You could say that the AI algorithm is the interface, a Large Language Model you have a conversation with through ChatGPT. Or should we be more specific, and discuss a speech recognition model like Whisper, converting any language you speak to the computer into instructions that are comprehensible for that computer?

## But where does the interface start and where does the computer end?

Maybe we should view the voice as an extension of the person, as it reaches further than the physical self. The voice could in this sense be compared to a computer mouse, only connected to us rather than the computer. We therefore get more entangled with the device itself, not in the least because we have placed part of our own intelligence into the computer. We can talk with the devices after we fed them with our thoughts.

At moments like these, you might be inclined to call on a GPT and ask for more clarity. But for now, we decline any artificial help with this report and try to see how far we can get with our own brainpower. After all, 'AI as an interface' can intuitively be understood without too defined a definition. When we look at the market capitalisation of NVIDIA, the company that supplies hardware for the server farms that the new language models run on, there doesn't seem to be any doubt.



## AI is becoming the new interface.

In this report, we present this new interface as a phenomenon from three different perspectives: it is an artificially intelligent ability, it is an activity – prompting – and it is a new form factor (new devices running on LLMs). These three perspectives bring a provocative vision of the future in which the new interface actually becomes intelligent. We choose to view the developments primarily from the perspective of the user, the consumer, precisely because the user has become increasingly deterrent for the


internal and external IT agenda of organisations. Where the employee previously brought a laptop and smartphone to work and became the driving force behind their adoption, this now happens more and more with AI.

We call the first lens, that of the activity, the **prompt factor**. The prompt factor stands for the new conversational interaction behaviour with the computer through prompting. Conversing through a prompt replaces swiping, just as it has replaced surfing before, which in its turn replaced punching.

Then there is the intelligence itself, the ability (or inability, if you wish): the **ability factor**. Here we mainly see important developments in terms of software agents that on the one hand prevent AI from flying off the rails and on the other hand change the interface from a passive into an active and sometimes even an autonomous actor.

Then finally, there is the lens of the devices, the **form factor**. The LLM robots and LLM gadgets introduced to the market, form this physical interface. Devices like these are also known as ‘connectables’. OpenAI recently gave a demo of its ‘Artificial General Intelligence’ robot, causing quite the commotion in the world. The controversial rabbit r1 device with its so-called Language Action Model promises to fundamentally change the way we interact with computers.

But the most important question of this report is whether we really see it all.





## Does this form a bridge to a new information era, a new supercycle?

After all, errors in judgement are easily made and of all time. But if the market is right, there is little room for doubt. The AI saga around the firing of Open AI's frontman Sam Altman basically tells the whole story. There is so much belief in the rise of the intelligent interface that the pushback by OpenAI's board members is an understandable, natural reflex.



# Chapter 2: Errors in Judgement Are Timeless

They are unforgettable images, etched in the collective memory, Steve Ballmer, then CEO of Microsoft, who couldn't stop laughing when in 2007 he was asked about his thoughts on the newly released Apple iPhone.<sup>6</sup> "500 dollars? Fully subsidised? With a plan? That is the most expensive phone in the world," he flaunted. "And it doesn't appeal to business customers because it doesn't have a keyboard, which makes it not a very good email machine," Balmer continued, confident in himself. "It may sell very well or not. We have our strategy. We've got great Windows mobile devices in the market today. You can get a Motorola Q phone now for 99 dollars. It's a very capable machine, it'll do music, it'll do internet, it'll do email, it'll do instant messaging. I looked at that and I like our strategy," a confident Ballmer concluded.

In the years after, Ballmer's bravado became painfully clear due to the explosive growth of the iPhone's market share as the favourite of the consumer, who appeared to be not so devoted to the keyboard and quickly mastered 'swiping'.

---

<sup>6</sup>Duivesteyn, S. (2011, 6 July). Steve Ballmer Laughs At The iPhone [Video]. YouTube. <https://www.youtube.com/watch?v=qycUOENFIBs>



The collapse of Ballmer's smartphone strategy was marked in April 2014 with Microsoft's acquisition of Nokia's mobile phone division. For 7.2 billion dollars, Microsoft took over the design, production, and sales activities of Nokia's mobile devices. Just one year later, in 2015, Microsoft had to take the hit of a 7.6-billion-dollar loss as the failed investment depreciated. In 2016, Microsoft sold the division to then pull back from the smartphone market altogether.

Steve Ballmer's error in judgement isn't unique in the world of tech. More of the greats have missed the mark. In 1943, IBM president Thomas Watson spoke the now legendary words: "I think there's a world market for maybe five computers." Ken Olsen, the founder of Digital Equipment Corporation, similarly announced in 1977: "There is no reason for any individual to have a computer in their home." In 1998 renowned economist and Nobel Prize winner Paul Krugman chipped in with his statement: "By 2005, it will become clear that the internet's impact on the economy has been no greater than the fax machine's." And: "As the rate of technological change in computing slows, the number of jobs for IT specialists will decelerate, then actually turn down; ten years from now, the phrase information economy will sound silly."


These miscalculations can be attributed to the lack of early recognition of changes that would have a profound influence on the way people handle information. For example, Thomas Watson and Ken Olson didn't foresee that computers would evolve from large, complex systems into user-friendly devices for everyday use. Paul Krugman underestimated the influence of the development of user-friendly web browsers and online platforms on the way that internet communication, information dissemination, and business practice would transform. Steve Ballmer didn't consider the move to touch screens, mobile internet access, phone apps,



and social media – all elements that brought fundamental changes in how people communicate and process information. Insufficient understanding of new technologies, not being able to foresee new breakthroughs and applications, and underestimating the speed at which changes happen, often lead experts – just like ‘ordinary’ people – to have prejudices and limitations.

But for every blunder of Krugman, Watson, Olsen, and Ballmer, there are also numerous innovative finds that have been put to the grave prematurely. Take for instance the ‘radio hat’, a metal pith helmet mounted with two radio tubes, an antenna, and built-in headphones that was launched with great fanfare at the end of the 1940s. The ad campaign at the time was accompanied with photos of young people enthusiastically walking around the city wearing the radio hat. It is reminiscent of the current images of people going outside wearing Apple’s Vision Pro (the newest AR/VR visor) as if it is the most natural thing in the world. The mentioned staging hasn’t helped the makers of the radio hat, and it remains to be seen if Apple will achieve the desired success. Endless products have never had a taste of the success that they were believed to achieve at the time of their introduction. But today, experts as well as the general public have little doubt about what will be the next big thing: artificial intelligence. IT mastodon and founder of Microsoft Bill Gates phrased it like this:

“The development of AI is as fundamental as the creation of the microprocessor, the personal computer, the Internet, and the mobile phone. It will change the way people work, learn, travel, get health care, and communicate with each other. Entire industries will reorient around it. Businesses will distinguish themselves by how well they use it.”



Bill Gates' strong language can be read in his statement 'The Age of AI has begun'.<sup>7</sup> The figure on the next page of investment company Coatue summarizes Gates' story nicely.

The old computer chips plus software were the keys to the computer age. The new chips plus AI models are the keys to the smart age. From the calculating computer, we're moving to much more comprehensive brain-like computers.

With their investments in OpenAI, Microsoft seems to have secured a bigger win. If all predictions come true, the Nokia debacle will soon be compensated for. And we all profit in an economic sense. According to the experts, generative AI will not only bring change in work, care, and communication, but it will also bring about a growth impulse into the economy. Analysts state that a yearly amount of 2.6 to 4.4 trillion dollars could be added to the growing economic wealth.

### A Quite Revealing AI Saga

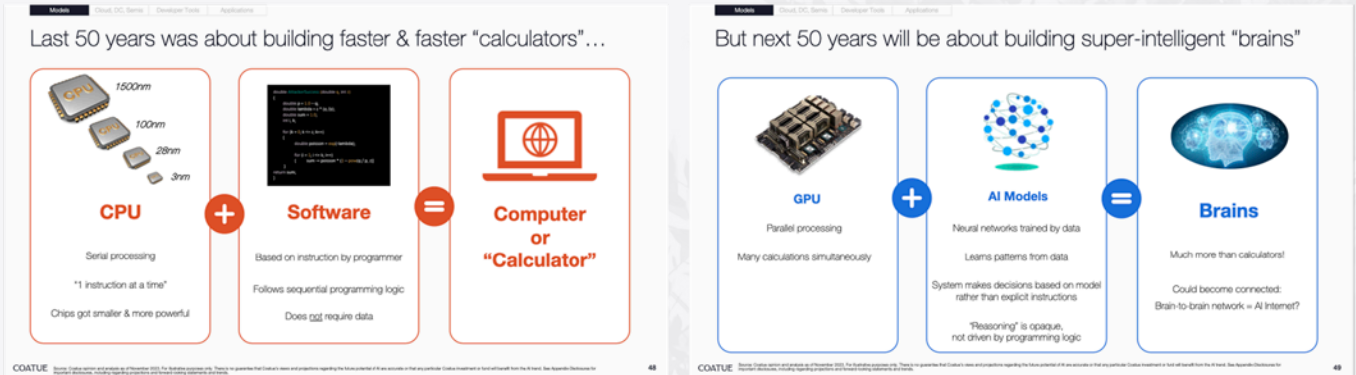
There is little doubt whether AI is the next big thing. Even more so, the conviction that this is going to happen is so omnipresent that here and there breaks are being put on in both word and deed. A peek behind the scenes at the makers of ChatGPT is a prime illustration. The company showed remarka-

ble restraint for their products to conquer the world. In a plot that would be fitting a Hollywood production, the world of AI was shaken by a true 'AI soap opera' within the walls of Open AI in November 2023. This saga, with OpenAI CEO Sam Altman in the lead, unfolded containing all ingredients of any good drama: sudden dismissals, regime changes at its top, and an internal conflict within one of the most influential companies in artificial intelligence.

Tensions reached boiling point when Altman, after a surprising dismissal, briefly took shelter at Microsoft. This dramatic turn led to unprecedented turmoil at OpenAI as well as the wider AI community. It even led to threats of mass departures of their 770-strong workforce. But like a phoenix rising from the ashes, Altman made a spectacular comeback as CEO, accompanied by a new board of directors, and promises of renewed priorities and better governance structures.

---

<sup>7</sup>Gates, B. (2023, 21 March). The Age of AI has begun. *GatesNotes*. <https://www.gatesnotes.com/The-Age-of-AI-Has-Begun>



Source: Coatue (2023, November). *The AI Revolution*. <https://drive.google.com/file/d/1gQhYT7j6b2wJmrFZHNeQgTiWPyTsjOfX/view>

# FROM INFORMATION TO INTELLIGENCE REVOLUTION

You may think that Altman was side-lined because of poor results as well as disappointing products and services but the opposite was true. ChatGPT was rushing into things, the general public was rushing into things. The board was afraid that their AI would derail from a lack of guardrails. Almost like Steve Jobs would have been fired on the eve of the original iPhone's launch because Apple's board of directors was afraid that the iPhone would derail.

In hindsight, you could say that the latter did happen. Socially and mentally the use of iPhones has changed us, and not always for the best. Only now we are building guardrails, like barring smartphones from the classroom, and seriously looking into rules around social media apps. This often happens, safety measures lagging behind the wish to the market. The dynamics behind this are outlined beautifully in a book that describes the advance of the car: *Unsafe at Any Speed: The Designed-In Dangers of the American Automobile*.<sup>8</sup> This 1965 book makes clear that, after a turbulent launch, technology will only be tamed in the subsequent years. In case of the automobile, these were mandatory safety belts, the roll cage,

<sup>8</sup> Nader, R. (1965). *Unsafe at Any Speed: The Designed-In Dangers of the American Automobile*. Grossman Publishers. For more information, see: [https://en.wikipedia.org/wiki/Unsafe\\_at\\_Any\\_Speed:\\_The\\_Designed-In\\_Dangers\\_of\\_the\\_American\\_Automobile#:~:text=Unsafe%20at%20Any%20Speed%20is,industry%20disregarding%20technically%20based%20criticism](https://en.wikipedia.org/wiki/Unsafe_at_Any_Speed:_The_Designed-In_Dangers_of_the_American_Automobile#:~:text=Unsafe%20at%20Any%20Speed%20is,industry%20disregarding%20technically%20based%20criticism)

and the construction of guardrails on motorways. The latter is what is currently at stake with AI in a metaphorical sense.

So far, the reasoning of OpenAI's board can be followed. Especially as commercial success wasn't originally the driving force of the company. Caution therefore prevailed. With the return of Altman, we can state that the big AI experiment started without restraint (see our previous report *The Great AI Experiment*). No holds barred! To what extent this move has been a dangerous one, only time will tell.

When Elon Musk got involved in OpenAI's saga (through his beloved platform X, of course) he asked if drastic measures like firing the CEO were necessary at all. And he added, "If OpenAI is doing something potentially dangerous to humanity, the world needs to know."<sup>9</sup> That there is bad blood between Musk and OpenAI only became apparent when Musk filed a case against OpenAI in March 2024.<sup>10</sup> As cofounder of OpenAI, he accused the company of deviating from its original course. It should be a not-for-profit organisation based on open source. Its current commercial course displeased him immensely. OpenAI responded by publishing their email traffic with Musk, from which emerged that in 2016 he agreed with the idea that OpenAI eventually would become less open and more commercial, for profit. The fact that Musk himself is marketing the competing xAI gives food for thought. It is clear that the book about the turbulent rise of OpenAI cannot be closed just yet.

---

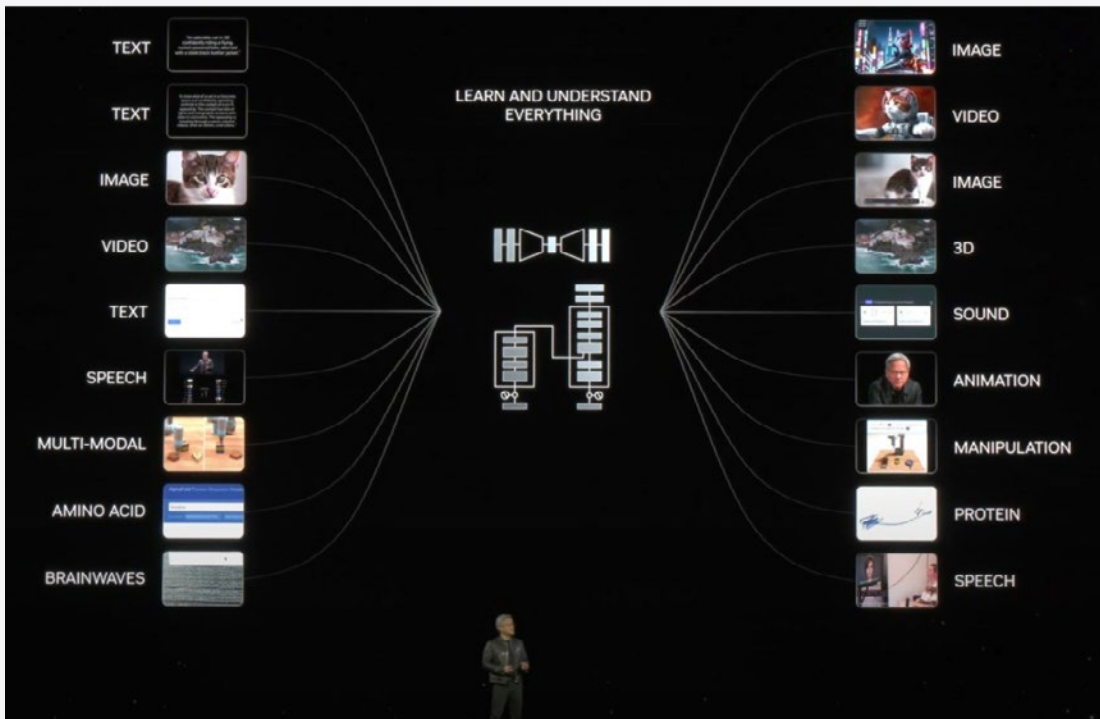
<sup>9</sup> Yao, D. (2023, 28 November). Why OpenAI Fired Its CEO Sam Altman. *AI Business*. <https://aibusiness.com/nlp/why-openai-fired-its-ceo-sam-altman>

<sup>10</sup> Verhagen, L. (2023, 7 March). Ruzie over toekomst AI escaleert: OpenAI publiceert interne mails van Elon Musk. *de Volkskrant*. <https://www.volkskrant.nl/tech/ruzie-over-toekomst-ai-escaleert-openai-publiceert-interne-mails-van-elon-musk~b0c2c796/?referrer=https://www.google.com/>



## The question if commercial brazenness leads to (un)expected AI dangers remains relevant.

The rumour mill keeps spinning, especially after Musk's insinuation that OpenAI has something to hide. Especially the secret OpenAI project 'Q' is fodder for far-reaching speculations. It is in essence a specialist AI programme that can solve mathematical problems. Experts say that this is a different kettle of fish to what OpenAI has shown up till now. Rumours claim that this would bring us one step closer to Artificial General Intelligence: an AI that matches or even surpasses humans in intelligence. Some equate the rise of AGI to the end of humanity.



'The intelligent interface' explained by Jensen Huang, NVIDIA's CEO, during a presentation for the general public.

Fortunately, Yann LeCun, head of AI at Meta and one of the founders of the field responds more realistically: “Please ignore the deluge of complete nonsense about Q\*. [...] Pretty much every top lab (FAIR, DeepMind, OpenAI etc) is working on that [...] It is likely that Q\* is OpenAI’s attempt at planning. They pretty much hired Noam Brown [...] to work on that.”<sup>11</sup> At a conference one day later, LeCun tells that we don’t have to be afraid that artificial intelligence will dominate humans. He calmly explains <sup>12</sup> that intelligence and dominance are two very different things. People want to dominate, but machines don’t. He adds that there is no question if AI will become smarter than humans. That will happen. In that near AI-intelligent future, every interaction we have with the world around us, will happen through an AI agent. And all those AI agents together are the sum of all human knowledge.

In this way, LeCun predicts that the interface of the future, the way we interpret the world around us and interact with it, will be an artificially intelligent system. We can assume that after MS-DOS, the web browser, and the smartphone, a new interface will arise, provocatively looking into a future in which artificial intelligence has permeated all capillaries of society. An interface between reality and perception; an artificial ‘transformer’ like the one that has been present in digital cameras and smartphones for a while. Every holiday snap is ‘upgraded’ to a slightly sharper and more attractive photo image thanks to AI. The upgrading gadget is a metaphor for future computers, changes to existing devices,

new AI-first devices we are not familiar with yet, new consumer products, and business software. We call the sum of these ‘the intelligent interface’. Jensen Huang, NVIDIA’s CEO, explains it as follows. All information that can be converted into a digital language, from text and images to amino acids and brainwaves (as shown on the left), can be used to recognise patterns. By systematically capturing those patterns, we can add meaning to them, enabling AI to learn and understand. With the recognised and understood patterns, we can then generate new language, images, and even proteins.

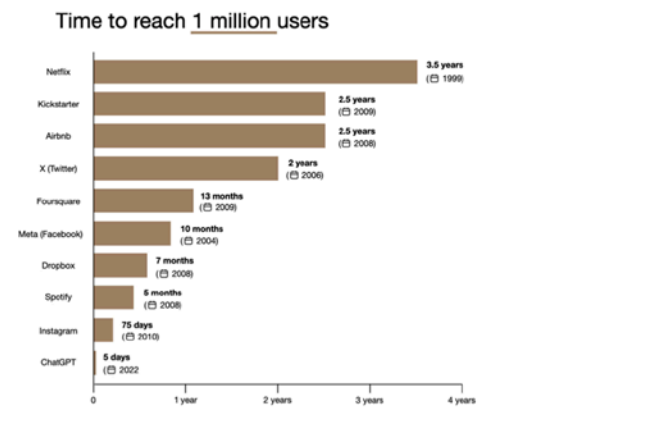
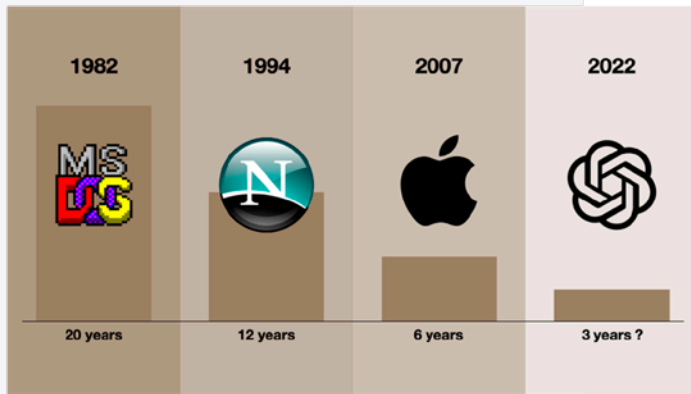
This creates an intelligent interface between all available information and the tasks we want to accomplish, situated at the heart of AI.<sup>13</sup>

<sup>11</sup> LeCun, Y. [@ylecun]. (2023, 24 November). *Please ignore the deluge of complete nonsense about Q\*. One of the main challenges to improve LLM reliability is to* [Post]. X. <https://twitter.com/ylecun/status/1728126868342145481>

<sup>12</sup> JJ [@JosephJacks\_]. (2023, 25 November). *Can we just make @ylecun president of AI and call it a day please?* [Post]. X. [https://twitter.com/JosephJacks\\_/status/1728510229133119644](https://twitter.com/JosephJacks_/status/1728510229133119644)

<sup>13</sup> NVIDIA. (2024, 18 March). *GTC March 2024 Keynote with NVIDIA CEO Jensen Huang* [Video]. YouTube. <https://www.youtube.com/watch?v=Y2F8yisiS6E>

# Chapter 3: The Prompt Factor



Altman feels empowered by the adoption numbers of OpenAI’s ChatGPT. He presented them, eleven days before his removal, on the stage of a developers’ conference. This presentation, that is reminiscent of Steve Jobs’ launch of the first iPhone, symbolises the transition into a new phase in technological development. Altman explains how in less than a year, generative AI has skyrocketed. Where Netflix needed roughly 3.5 years to have 1 million people commit, text generator ChatGPT reached this milestone in only 5 days.

If we list ‘the big four’ companies that fundamentally changed the handling of computers and information, an image of ever shorter innovation cycles emerges: each innovation reaches broad acceptance (80% of a population) in half the time. MS-DOS took 20 years, Netscape 12 years, the smartphone 6 years, and OpenAI is estimated at just 3 years.

Not just the handling of information changes, the size of the equipment used changes as well. The mainframe computer used to take up an entire basement, and later on the personal computer conquered the desk. Subsequently the laptop got a hold of our lap, to finally end up as a smartphone in our trouser pocket or handbag. We even take the

device to the toilet and save a special place for it on our nightstand. Thanks to artificial intelligence, the next generation of computers won’t even need a screen anymore. New interfaces based on NeuroLinguistic Programming (or NLP), and Natural Language Understanding (NLU) make it possible to have a conversation with the device and input prompts through speech. When our interaction with computers becomes this straightforward, we’ll forget the background technology and only magic remains. Or, as science fiction author Arthur C. Clarke once said:


“Any sufficiently advanced technology is indistinguishable from magic.”




## Punching, Surfing, Clicking, Swiping, and Prompting

If we travel through the history of computer communication with giant strides, we'll find that the intermediary layer in between human and computer or interface has drastically changed in the past ninety years. As a starting point, we'll take the punch card, that came into use in the 1930s. On specially designed machines, information was perforated into paper cards, after which those cards could be read by the computer. The information on the punched cards was indecipherable by humans but much more understandable to the computer.

We communicated via a code language to get the computer to run. With the rise of the monitor and keyboard, the coding on punch cards could be omitted and instructions could be given in a more direct manner. But it took until the introduction of MS-DOS for the computer to become available to a wider audience. Until then, computer use was limited to companies and universities, and it was hard to imagine to ever own a personal computer. It was the simplified interface offered by MS-DOS that brought about the PC revolution in the late eighties, and finally made the computer commonplace. The introduction of the graphical interface of Apple OS initially, and the introduction of Windows 1.0 in 1985 further incentivised the PC revolution.



A punch card and a card puncher. A punch card's capacity usually was 80 bytes.<sup>14</sup>



<sup>14</sup> Photos: <https://twobithistory.org/2018/06/23/ibm-029-card-punch.html>

A similar phenomenon was seen with respect to the internet. Although the official birthday of the internet was set at 1 January 1983, it took more than ten years before it resonated with the general public. Only people working at universities and research organisations had access, and with rudimentary applications such as Gopher and Archie the information could be consumed. With the introduction of a user-friendly web browser by Netscape, easy access to hyperlinked information of the growing world wide web was realised. The simplified interface through which people ingested information triggered another revolution, the internet revolution.

A next innovation of the interface through which people take in information, was presented in 2007 with the introduction of the iPhone by Apple's Steve Jobs.

Initially presented as a three-in-one device – a music player, an email device and an internet browsing device – the smartphone was now a reality.

The resulting 'apps' changed the way people consumed information again, and another revolution took place. It is more than likely that with the introduction of ChatGPT and the prompt interface a revolution will be triggered, although we can only determine the essence of it in hindsight. Yesterday's interface of punching has successively given way for typing, clicking, surfing, and eventually prompting in natural language.



## MS-DOS: The Command Interface

MS-DOS is a command interface, a straightforward interpretation programme that allowed you to handle the computer. Above all, it is a communication standard that allowed programmers to develop software that ran on this interface. Because PC builders accepted this standard, it was possible to build software on a large scale that could then be installed on every PC. That such an interface takes twenty years before it is widely used, isn't surprising. With such an MS-DOS machine, you could perform some chores, but you couldn't easily go online with it, or order things on Amazon. That came much later. The MS-DOS interface made it possible to mix business and pleasure by accessing other programmes on a personal computer that were useful or interesting to you, including the erotic Leisure Suit Larry games (with the last edition being released in 2020 and running on Windows, among other platforms). Until then, computers were mainly found in the basements of large organisations to perform rigid maths.

## The Graphic User Interface: Windows and Apple OS

Reportedly, it was Steve Jobs who during a visit to the Xerox Parc research institute in Palo Alto was demonstrated a rudimentary version of what we now know as a graphic user interface and immediately understood its value. No longer was it necessary to memorise complicated instructions to load application software onto a PC. A simple graphic presentation of a desktop, an application window, icons, menus, and a mouse would dramatically simplify the interface. And it was Apple who in 1984 tried to counterbalance the dominance of the IBM compatible MS-DOS PC with the introduction of the Macintosh computer. Although Microsoft tried to offer a comparable interface as layer on top of MS-DOS in 1985 with Windows 1.0, it would take years before this WIMP interface became commonplace with the introduction of Windows 3.0 in 1990. The command interface disappears into the background and all information on the PC was able to be unlocked with a few simple 'clicks'.

## Surfing


The word alone makes it clear how different the use of computers became with the introduction of the web browser. On the waves of data, we let ourselves drift away on an ocean of information and possibilities. Purposeful clicking is starting to get competition from meandering surfing. In the early years, you would open the Internet Yellow Pages to find out where to surf to. In this physical Yellow Pages of the internet, you could find all sites of the world wide web. Search engines took over


when the internet started to burst at the seams. Nowadays, you can run ChatGPT within your search engine to evade its wave of commerce (the first hits invariably are sponsored items). You don't just navigate to information now; the web browser became a thing to use. You use it to plan your holidays, birthday gifts, your taxes, and your pension. E-business and e-commerce received a boost due to the surfing capabilities of browsers.

### Swiping

Engaging with computers became a much more sensory experience with the arrival of the iPhone. It's not just the swiping and camera with face recognition for verification to unlock apps, it literally became more sensory. With the wind blowing in your hair, on the playground, on a boat, or in the doctor's waiting room, in any circumstances a swipe toward information, a booking, or a TikTok video can be made. Because of the smartphone, computer interaction became the 21st century lifestyle. The fight for our attention definitively changed – are you on your phone or in the here and now? The 2010 classic *Never Look Up* shows photos of people on the streets that are all looking down at their phones. And the book *Alone Together* by Sherry Turkle (2011) seriously questioned our obsessive information behaviour.

This much is clear: the smartphone has changed us as people.





## Prompting

To fully reduce the AI revolution to the textbox in which we formulate our assignments (prompting) might be giving a bit too much credit to the action itself. But it is a good starting point: punching, typing, clicking, surfing, swiping, and prompting. Prompting immediately classifies itself as the mother of all command structures. From the code language of the punch card to the natural language that parents and teachers brought us up with. The power of prompting is hidden in the dialogue. By questioning AI step by step, you get much better results sooner than you do with search queries. Just like people engage in conversation to get to the core of the matter, conversation is the essence of prompting. Which explains why prompt engineers are currently worth so much money.

They know like no other how to work AI to get better results.

First and foremost, prompting contributes to the idea that we can easily set the parameters ourselves. You want to see a photo that looks so and so, to make a computer programme that allows a red ball bounce, or to write a text in the style of *The Daily Telegraph*. Your own wish, worded in your own language, forms the command for the device. Other skills or translations aren't necessary anymore. Type or speak your wish, and the magic can begin. The deeply human desire to be in control is satisfied in this way.

## The Prompt Economy Needs Chips

To enforce the idea of AI as the mother of all interfaces, it would be good to take note of the unbridled ambition of Sam Altman. He is on a tour to collect 7 trillion (7,000 billion) dollars in the market for OpenAI's new plans.<sup>15</sup> Maybe that sum of money includes the 100 billion dollars to build supercomputer Stargate with Microsoft. The rumours that this is going to happen are becoming stronger.<sup>16</sup> To satisfy his appetite for money, Altman is visiting the United Arab Emirates. But you read it right: 7,000 billion dollars. And all that to develop the AI computer chips that are necessary to feed the whole world with the – in his eyes – much-needed AI support. To realise OpenAI's AI interface strategy, more powerful chips are simply necessary.<sup>17</sup> The fact that the amount is this astronomical (about 8% of the world's economy), makes it clear that Altman's ambition knows no bounds. For a critical note on the scarcity of resources such as computer chips in a world with finite resources, we refer to our latest report, *The Rise of the Experience Ecology*.<sup>18</sup>

The fact that Altman has his sights set on AI chips isn't surprising.<sup>19</sup> If anyone is benefiting of the new interface becoming intelligent, it will be chip manu-



facturers. That is what NVIDIA's numbers are showing. In 2023, the financial world was amazed by the company reaching no less than one trillion dollars in market value. NVIDIA's spectacular growth is illustrated by a stock price rise of 726% (from 108.13 dollars in 2022 to 785.75 dollars at the end of February 2024). For every dollar turnover, the company by now makes a gross profit of 76 cents.<sup>20</sup> This can all be attributed to the insatiable demand for GPUs (Graphical Processing Units), the crucial building blocks of AI that NVIDIA's founder Jensen Huang invested in early on, at the time as supplier of particularly the video game industry.

The select group of biggest companies of the world, measured by their market value, gives an interesting image of the changes in the economy and technology throughout the years. Previously, oil companies such as ExxonMobil and Shell topped the rankings. This shows the economic interest of oil at the time. But with the arrival of computers, the internet, and mobile phones, IT companies such as Microsoft and Apple took over their position as most valuable companies.

---

<sup>15</sup> Hagey, K. (2024, 8 February). Sam Altman Seeks Trillions of Dollars to Reshape Business of Chips and AI. *The Wall Street Journal*. <https://www.wsj.com/tech/ai/sam-altman-seeks-trillions-of-dollars-to-reshape-business-of-chips-and-ai-89ab3db0>

<sup>16</sup> Gardizy, A. & Efrati, A. (2024, 29 May). Microsoft and OpenAI Plot \$100 Billion Stargate AI Supercomputer. *The Information*. <https://www.theinformation.com/articles/microsoft-and-openai-plot-100-billion-stargate-ai-supercomputer>

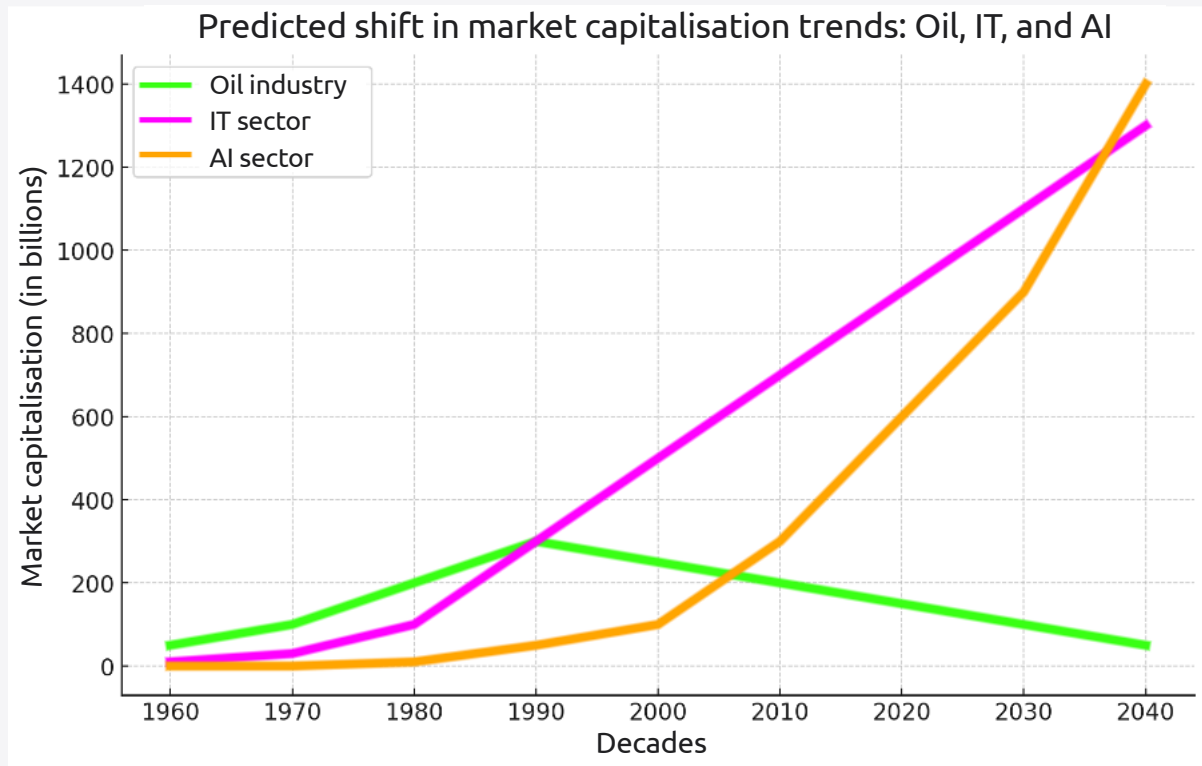
<sup>17</sup> Al Jazeera (2024, 9 February). OpenAI's Sam Altman seeking trillions to fund chips for AI, report says. <https://www.aljazeera.com/economy/2024/2/9/openais-sam-altman-seeking-trillions-to-fund-chips-for-ai-report-says#:~:text=OpenAI%20CEO%20Sam%20Altman%20is,Wall%20Street%20Journal%20has%20reported>

<sup>18</sup> See <https://ict-books.com/book/the-rise-of-the-experience-ecology/>.

---

<sup>19</sup> Rumour has it that Meta and Amazon are venturing into the chip market. These organisations are wholesale buyers of computer chips.

<sup>20</sup> Vogels, P. (2024, 22 February). Gekte bereikt nieuw hoogtepunt: AI-chipfabrikant Nvidia in één klap 200 miljard meer waard. *Algemeen Dagblad*. <https://www.ad.nl/economie/gekte-bereikt-nieuw-hoogtepunt-ai-chipfabrikant-nvidia-in-een-klap-200-miljard-meer-waard~a908bd80/?referrer=https%3A%2F%2Fwww.google.com%2F>



Historical stock quotes of the top ten largest companies in market value further extrapolated by GPT-4.

Now we see another shift, this time towards companies that specialise in AI, like NVIDIA, which points to a new phase in the economic development in which AI plays a crucial role.

The chart shows that the market capitalisation of AI surpassed that of IT. It is also a call to all stakeholders (us all) to prepare for a future in which AI is not only part of the economy but will be a fundamental motivation of it. If we dare to believe GPT-4, as we asked it for further extrapolation, the AI sector will surpass the IT sector in market value halfway through the next decade. The question remains if you can make a distinction between the two sectors by then.



## Chapter 4: The Capability Factor



The most important question of course is what the ‘capability’ is we add to our digital existence by using an intelligent interface. The easiest answer is: artificial intelligence. The one who coined the phrase, John McCarthy, originally somewhat reluctantly came up with the phrase in 1956. He really found it too ‘flashy’, a bit of a gaudy phrase. But he also had bad experiences with other labels for the new discipline. His earlier attempt to call the field ‘Automata Studies’<sup>21</sup> floundered. The name that Nobel Prize winner Herbert Simon came up with, ‘Complex Information Processing’, didn’t make it. Researchers in this field didn’t believe each other’s information systems were complex enough and bickering dominated the discussion. So, the phrase ‘artificial intelligence’ came about rather reluctantly. McCarthy confessed that it was merely a cunning marketing move:

“Excuse me, I invented the term ‘Artificial Intelligence’. I invented it because we had to do something when we were trying to get money for a summer study in 1956, and I had a previous bad experience.”


John McCarthy

---

<sup>21</sup>‘Automata’ is the plural of ‘automaton’, in independently operation machine (see [https://nl.wikipedia.org/wiki/Automaton\\_\(machine\)#:~:text=Een%20automaton%20\(meervoud%3A%20automata%20of,reageren%20op%20vooraf%20bepaalde%20instructies\)](https://nl.wikipedia.org/wiki/Automaton_(machine)#:~:text=Een%20automaton%20(meervoud%3A%20automata%20of,reageren%20op%20vooraf%20bepaalde%20instructies).)).

## Capability and Incapability

Intelligence is a concept from cognitive science, which is one of the humanities. But people aren't computers, and computers aren't people. But this 'cognitive borrowing' is very usual in the emergence of a new field, simply because words you could give to the developments don't exist yet. This is not just true for computer science, but for neuroscience as well. The latter strikingly enough consulted computer science to find the right wording. Simply put, they started looking for human terminology for new developments in computer science and computer terminology for developments in humanities. The dangerous side effect is that we will see brains (or humans) as calculators in which information can be stored and processed, and computers as a psychological phenomenon that will for instance have a capacity for hallucinating. But human hallucination and that of a computer don't have anything in common. We see things that aren't there and they 'show things' that you wouldn't expect based on the input of the computer systems. Take the concept 'attention', for instance. It is used in the field of artificial intelligence as well as in cognitive science, but the meaning is very different in each specific context.



### Artificial Intelligence

"Attention is a mechanism, within neural networks, particularly transformer-based models, that calculates "soft" weights for each word, more precisely for its embedding, in the context window."

Wikipedia

### Cognitive Science

"Attention is the concentration of awareness on some phenomenon to the exclusion of other stimuli. It is a process of selectively concentrating on a discrete aspect of information, whether considered subjective or objective."

Wikipedia

The problem that this type of language theft creates is described in detail in *Anthropomorphising machines and computerising minds: the crosswiring of languages between Artificial Intelligence and Brain & Cognitive Sciences*.<sup>22</sup> This report by information philosopher Luciano Floridi and neuroscientist Kia Nobre raises the alarm for further development of the field. According to Floridi and Nobre both will keep the wording but the meaning we give to artificial intelligence will receive an 'upgrade'. For instance, no-one in this day and age actually thinks that the sun really rises at sunrise. We have known long enough that the earth revolves around the sun, and not the other way around. We take our coffee with sugar or a sweetener and know that they are two different things. When we're talking about horse power under the hood, no-one is thinking about horses anymore. There will be a day that we no longer think about intelligence when talking about AI, but about capability, like with 'horse power', and the incapability of that which AI can't or isn't.

There are also calls to name artificial intelligence a 'new digital species' from now on. Or that is the position of Mustafa Suleyman, Microsoft AI's CEO. He is also of the opinion, like Floridi and Nobre, that we aren't using the correct language and metaphors to explain AI.

According to Suleyman, we need new wording, or we cannot oversee and regulate AI.

---

<sup>22</sup> Floridi, L. & Nobre, A.C. (2024, 25 February). *Anthropomorphising machines and computerising minds: the crosswiring of languages between Artificial Intelligence and Brain & Cognitive Sciences*. Centre for Digital Ethics (CEDE). [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=4738331](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4738331)

## Tim versus the Police Algorithm

A clear example of a pitfall is shared by Tim Hanssen.<sup>23</sup> He received a 380-euro fine for using a phone in his car. The photo it was based on, indeed shows Tim with his hand by his ear but without a phone. The irony is that Tim himself works on AI image recognition and he was able to put the finger on the sore spot. According to him, something went wrong in the training set of the ‘intelligent’ system. It was tested with thousands of photos of people holding a phone to their ear. But just as many photos of people with just their hand by their head – without a phone – should have been taken. If that had been the case, the system could have easily spotted it. As a human reviewer is always involved to conduct a final check, in Tim’s case the same mistake was made twice. That person also assessed that Tim had a phone in his hand. Tim’s hope is that if you better explain the capabilities of the system to the people who check these results, they will pay more attention while doing these checks.




Tim Hanssen shows how a ‘false negative’ is created when training data aren’t in order.


## Woke by Design

Announced with much bravado, Google’s AI show-piece Gemini provides us with another example. Shortly after its introduction, the picture-making component of Gemini completely flew off the rails resulting in it being taken offline indefinitely. After a request to generate either an image of the pope or a Nazi officer from World War II, it invariably returned with people of colour. Google had to follow suit and acknowledge that this was never the intention of the model they designed.

Immediately, the phrase ‘woke by design’ did the rounds online as a possible explanation but Google is keeping quiet about the exact cause. It’s fair to say that every capability has its own limitations. In that sense, ‘capability’ is a better term as it gives an invitation to think about what exactly you can find under the hood of the newest machines.

<sup>23</sup>Hanssen, T. (2024). Tim versus het politie algoritme. <https://nippur.nl/tim-versus-politie-algoritme/>





Results of prompts asking for an image of the pope or of Nazis on Google's Gemini platform.

## Is History Repeating?

The suggestion that the computer is actually intelligent in the classical sense, is not confirmed by either example. We consider the software conference that was organised by NATO in Garmisch-Partenkirchen in 1967.<sup>24</sup> Its goal was to nip the anticipated software crisis in the bud. Software projects ran over by default, were above budget, and underperformed when measured by the code's quality. Incidentally, the term 'software engineer' is a result of that same 1967 conference.

At the time, it was still possible to gather the IT industry, as the worldwide number of software engineers was limited. The trigger of the conference was the software crisis. But the fact that the software code, once produced and distributed, would spread over the world, also raised concerns. Without the possibility of a call-back, incorrect code would automatically be replicated. That makes an error difficult to restore. To avoid similar problems, expertise was paramount and universal rules for software engineering needed to be implemented.

A repeat of that NATO conference, this time focused on AI, would not be out of place, especially in view of Google's recent emergency intervention. Fortunately, there are developments in the field of generative AI that could raise its quality making its outcomes more reliable.

<sup>24</sup> Naur, P. & Randell, B. (eds.). (1969, January). *Software Engineering*. <https://www.scrummanager.com/files/nato1968e.pdf>

# Not the human but the computer offers a helping hand in this case, and in the shape of new types of software agents.

## New Agent Models Come to the Rescue

Andrew Ng is a big fish among AI experts. In his blogpost of the end of March 2024, 'How Agents can improve LLM performance', he explains how software agents can greatly boost the results of an LLM.<sup>25</sup> Ng is the founder of DeepLearning.AI, among others, founder of learning platform Coursera and Adjunct Professor at the Computer Science department of Stanford University. He starts his message in his blog as follows:

“Dear friends,

I think AI agent workflows will drive massive AI progress this year — perhaps even more than the next generation of foundation models. This is an important trend, and I urge everyone who works in AI to pay attention to it.”

Andrew Ng

## Think Before You Speak

The principle of these new generation agent models (or 'foundation' models) is that rather than an instant ('one-shot') solution, a step-by-step model is built in. In this step-by-step model, the language model takes trips to other websites or reflects on

the outcome it just generated. Instead of a one-shot solution, you build a multimodal workflow in which agents within the system independently get started to then get the best results together. More simply put, let the LLMs think before they speak. And you can leave the thinking to the system itself, so we read in Stanford University research paper 'Quiet-Star: Language Models Can Teach Themselves to Think Before Speaking'.<sup>26</sup>

So, you're giving LLMs additional assignments, for example:

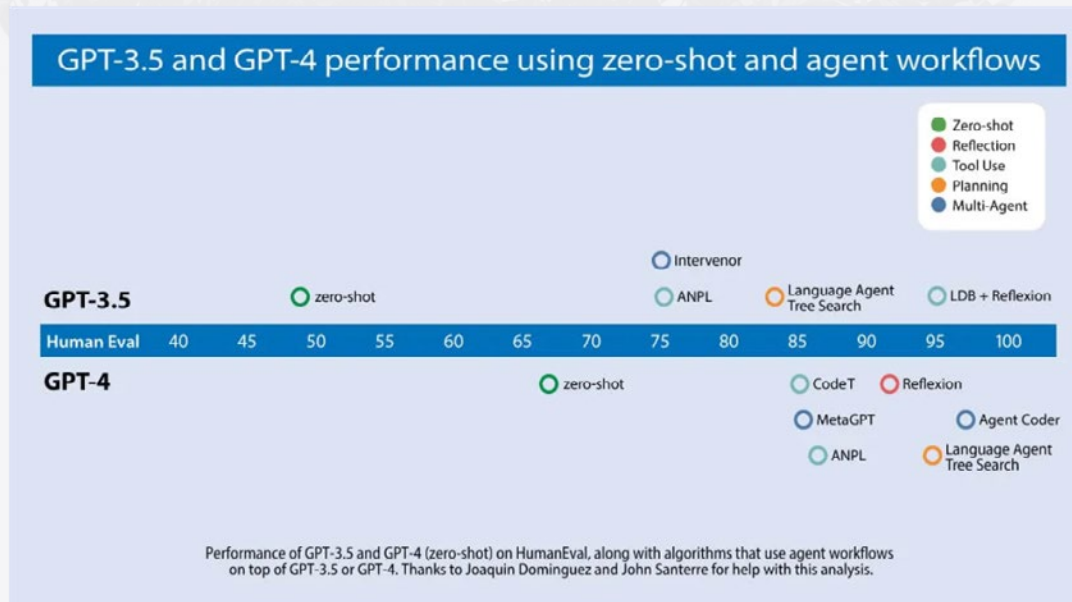
- Give an outline first.
- Then decide whether it is necessary to look for extra information on the web.
- Decide what you are going to do with it.
- Look at the initial outcome and look for unexplained outcomes.
- Adjust the first draft on the base of weaknesses that are found.
- Only then give an answer.

The agent models that Ng and his team have studied, pertain to the development of software. In this, we must realise that this is only a metaphor for each form of 'intelligence' and reasoning. Agents can also organise things for us, like pick out the best holiday and give advice on what to do, find the best tickets, the most affordable travel insurance, and so on.

---

<sup>25</sup> Ngu, A. (2024, 20 March). Agentic Design Patterns Part 1. *The Batch*. <https://www.deeplearning.ai/the-batch/how-agents-can-improve-llm-performance/>

<sup>26</sup> Zelikman, E. et al. (2-24, 18 March). Quiet-STaR: Language Models Can Teach Themselves to Think Before Speaking. arXiv. <https://arxiv.org/abs/2403.09629>



Source: The Batch (2024, 20 March). Robots Talk Back, AI Security Risks, Political Deepfakes, and more. <https://www.deeplearning.ai/the-batch/issue-241/>

The figure above shows how GPT-3.5 and GPT-4 score compared to such an agent model. These workflows, such as Large Agent Tree Search, MetaGPT, and Agent Coder, offer substantially higher quality. One of the lessons to be learned is that the GPT models don't have to keep getting larger to improve quality. Another important lesson is that major quality hikes are being made with such an agent model, and concerns of language models derailing diminish because of it.

## Agent is the Magic Word

### GPT Assistants

At the moment, many agent-based AI assistants are in the pipeline. Take OpenAI's GPTs as an example. These AI apps were launched at the end of 2023 as a new innovation. They are ready-made applications, but you could build your own GPT. An example of such a standard GPT is 'Sell me this pen'. It is a sales aid that will get to work for you. After uploading a photo of a product you'd like to get rid of (a bank, a shirt, or a road bike, for example), this seller will get to work for you. Firstly, the GPT will scour sales sites looking for products that are like the one you are offering for sale. Subsequently, the GPT will give a price indication within a range of a minimum and maximum amount. If there are any ambiguities, the GPT will indicate this (is it a silk shirt, or is it made of synthetic wool?). Next, an advertising text will be generated to sell the goods. (This is where language models are at their best.) The follow-up step is to place the ad, including a photo, on a platform like Whoppa or Vinted. If this is automated, you could be putting your entire wardrobe in the shop window within two minutes. It is a dream within arm's reach: GTPs managing GPTs.

“As intelligence gets integrated everywhere, we will all have superpowers on demand.”

With these words, Altman ended his visionary keynote on Monday 6 November 2023 on a developer conference in San Francisco.<sup>27</sup> He outlined a near future in which ChatGPT users can develop their own personalised digital assistants. Previously, he explained that these specialised applications, named GPTs (or ‘agents’), not only are able to read, write, hear, speak, and see, but are also able to create art, reflect, use existing computer equipment as a tool, specialise in certain topics, utilise specific data, execute actions in the digital world, and communicate and cooperate in a specialised way.

These GPTs can even take on different roles, varying from writing coach to sous-chef, from Maths teacher to creative design partner. Remarkably, creating these specialised GPTs doesn’t require in-depth programming knowledge. A new digital marketplace, comparable to Apple’s App Store, promised to further boost the development of these personalised assistants, fundamentally changing the way we interact with technology. Agents are the new apps, and prompts are the new queries.

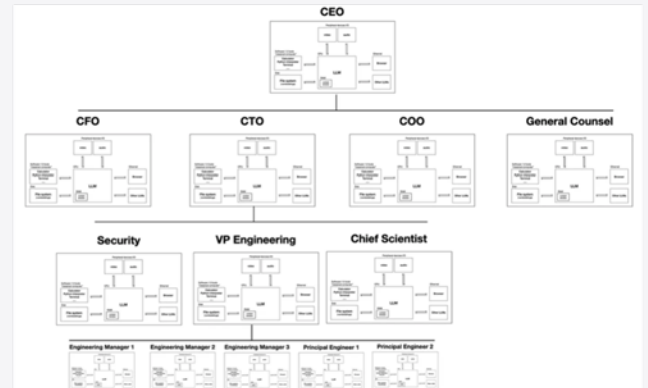
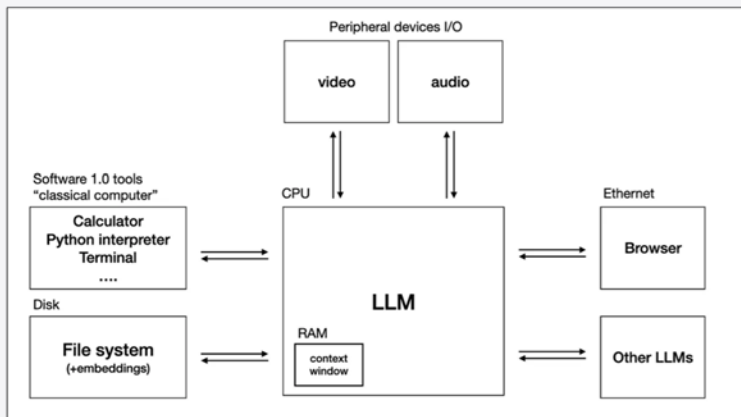


Imagine you’d like to invite a co-worker to a lunch at a specific location. Usually, this would take three separate smartphone apps: a diary to invite your co-worker, a platform to book a restaurant, and a travel app to plan the journey by public transport. In the vision of LLM as an operating system, you could simply say: “I’d like to have lunch with co-worker X tomorrow at 1pm at a location close to Utrecht Central Station.” It will translate your spoken instructions to the required tasks and will take action to achieve the desired result without the need for a range of apps. So, the power does not just lie in the interaction of user and LLM, but additionally in that which happens in the background, the interaction between LLM and other APIs, LLMs, and apps.

<sup>27</sup> Altman, S. (2023, 6 November). OpenAI DevDay: Opening Keynote [Video]. YouTube. <https://www.youtube.com/live/U9mJuUkhUzk?si=ao7ukQnQNbK70Ab>



### LLM as an Operating System



Source: Karpathy, A. [@karpathy]. (2023, 11 November). *LLM OS. Bear with me I'm still cooking.* [Post]. X. <https://twitter.com/karpathy/status/1723140519554105733>

### Healthcare Practice on Steroids

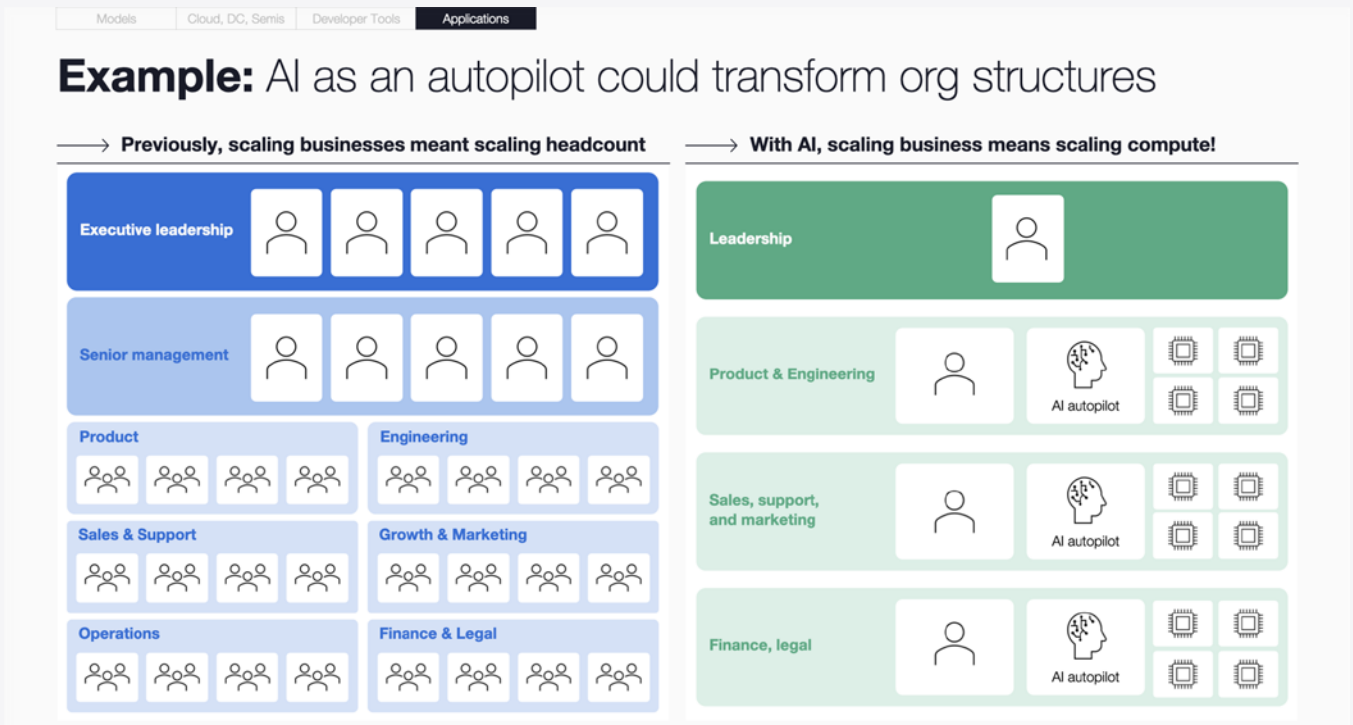
One practical example of how this agent technology could work in real life, is how psychiatrist Joan Mundin has organised her therapy practice with the aid of ChatGPT. In her waiting room sits her patient Robert Scoble, a technology analyst and tech author with a troubled childhood. During her initial consultation with Scoble, she runs her personalised agent system in the background. The system records the conversation (image and sound) and describes some characteristics, such as Scoble’s outfit. The objective transcript of their conversation serves as input for a subjective interpretation by the agent of what might be the underlying issue. Next, the agent generates the corresponding treatment plan. For the financial settlement of the plan, the agent will use the correct clinical coding system and generates the bill to be sent. According to the doctor, the quality of this software agent is comparable to that of a newly qualified student. Double-checking will always be needed and desired but in professions that have sworn the Hippocratic oath, that is usually

feasible. We know of this case, as Scoble himself went public with it,<sup>28</sup> despite the – at times – very personal details of his difficult childhood.

### AI as the New Operating System for Organisations

When we stretch the phrase ‘capability’ a bit further, we can also observe how AI can change organisational skills. For this, we consult Andrej Karpathy, Tesla’s former director of AI who is currently building a type of JARVIS (Just a Rather Very Intelligent System) for OpenAI (with the name derived from the AI system of Tony Stark from the Marvel franchise *Iron Man*). Clearly inspired by science fiction, Karpathy goes one step further. He sees a future in which language models operate as an independent operating system for organisations.

<sup>28</sup> Scoble, R. [@Scobleizer]. *Second, here’s a video of me talking with Dr. Mundin after our session where we go into what the notes* [Post]. X. <https://twitter.com/Scobleizer/status/1657808538847313920>



Source: Coatue (2023, November). *The AI Revolution*. <https://drive.google.com/file/d/1gQhYT7j6b2wJmrFZHNeQgTiWPYtjOfX/view>

The multimodal LLM is the central part that supervises the processing of incoming and outgoing information (in various formats), the translation of computer code, and the storage of specific data. In a way, the AI will have human senses such as hearing and vision. This enables artificial intelligence to interpret audio and video, to recognise patterns in it, and to then generate new audio and new video. The translation from the LLM as an operating system to instructions for different functional areas in an organisation has yet to be made. Every department – marketing and finance, for instance – has its own operating system, composed of specific language models with their own senses. Karpathy gives a fairly clinical view of an organisation without being very specific about how decision-making processes between departments would ultimately take place, how overall governance would be arranged, and how guarantees will be put into place to secure quality.

According to David Friedberg, founder and CEO of The Climate Corporation and co-host of the All-In podcast, in this, the classic client-server model is reversed.<sup>29</sup> Nowadays, we are both the client, an endpoint in the network, and we communicate with the centre of the network where information is gathered, and transactions are processed. In the future, the opposite will happen. When more and more of our data is gathered on our individual client, each one of us becomes a server.

It is an interesting and once more provocative thought that everything could be turned on its head when a language model evolves into an operating system. Positions within the organisation, as drawn out in the model above, will fully rely on the possibilities that LLMs offer.

<sup>29</sup> Douzet, A. (2023, 20 July). AI And The Race Towards Intimacy. *Forbes*. <https://www.forbes.com/sites/forbesbusinesscouncil/2023/07/20/ai-and-the-race-towards-intimacy/>

The above further implies that we could design our organisations, economy, and society in a completely different way. Karpathy imagines a world in which every position in an organisation, from CEO to engineering manager, would use a specialised LLM to optimise both decision-making and productivity. Routine tasks would be automated with unprecedented accuracy and complex problems would be solved by the combined intelligence of multiple LLMs in cooperation. Investment firm Coatue has the exact same thoughts.<sup>30</sup>

Coatue also see a future in which every employee in a company has one or more AI co-workers.


And whichever approach you take, the conclusion seems to be that true efficiency and quality improvements are found in the so-called capabilities of the agents. If they actually perform tasks in the background with sufficient efficiency and accuracy, we'll see an unprecedented acceleration within organisations.

This is a vision of the future with far-reaching consequences for the total number of employees within an organisation. It is the first time that the work of knowledge workers is affected by automation technology in such a way. According to a scientific paper, the introduction of generative AI was impactful in the short term:<sup>31</sup> shortly after the introduction of ChatGPT, the number of freelance copywriters decreased drastically.

---

<sup>30</sup>Viswanath, S., Khanna, V., & Liang, Y. (2023, 16 November). AI: The Coming Revolution. <https://www.coatue.com/blog/perspective/ai-the-coming-revolution-2023>

<sup>31</sup>Hui, X., Reshef, O., & Zhou, L. (2023, 31 July). The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market. [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=4527336](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4527336)





Chapter 5:  
**The Form Factor:  
The Next Big Thing**

The computer, laptop, and mobile phone are inherently devices we can reach our server parks with via prompt.

Within the foreseeable future, glasses, necklaces, pins, and other yet undefinable new devices will be added. Suddenly, our smartphone is no longer untouchable, and we may put this previously indispensable tool to the side. All new devices will have a key asset: AI as the new operating system, understanding of what is necessary, and the autonomy to complete tasks. This AI revolution poses the question what hardware will be necessary to carry this 'no-app' revolution. Let's investigate a number of these new devices and their background stories. Sam Altman is now collaborating with Jony Ive, former Chief Design Officer of Apple, responsible for iconic products such as the iMac, iPod, and iPhone, on an – in their own words – revolutionary product. Financed by a 1-billion-dollar fund of the Japanese SoftBank Group, they are trying to build the next form factor, the AIPhone, in which ChatGPT should operate as the device's brain. Finding out what Ive is going to come up with is a matter of time. But Altman is betting on several horses at once. He also invested in the AI Pin by American tech start-up Humane.





## The Ai Pin

The Ai Pin was released on 9 November 2023. It is an intelligent pin that wants to free the world of its smartphone addiction. Under their motto “A future with AI is better than a future with more screens”, the Ai Pin was announced with an excellent pitch at a TEDx conference.<sup>32</sup> Former Apple employees appear to have a hand in this development too. The founders of Humane, Imran Chaudhri and Bethany Bongiorno, worked together on product development at Apple for years. They are also life partners.

During his TEDx presentation, Imran lets his wife dial in and incidentally proves that you don't need a screen to be able to call. Which we knew from the

past, of course. The device projects text and monochrome images in the owner's palm so he can tell who's calling. When Chaudhri addresses his audience with his Ai Pin in different languages a little later, he adds that it is his voice they are hearing, and his own emotions. It is an AI clone of Chaudhri's own voice they are hearing, but in a different language. Communication with the Ai Pin runs through touch, hand gestures, and voice commands. The most interesting is the underlying operating system, named Cosmos, that doesn't function as a collection of apps, but as a system that can summon different AIs and other tools when you need them. A YouTube video shows how something like that works.<sup>33</sup> A man presses the Ai Pin with his finger and asks it to describe the woman who is sitting right in front of him. After a few seconds, a female voice answers and the lady in question is perfectly described: The woman sitting in the chair is wearing dark sunglasses and a camouflage-patterned jacket with a black dress or skirt underneath. She is holding a small yellow object, possibly a toy. It is not a toy, but an apple. The fact that the Ai Pin wasn't quite sure could be inferred from the subtle addition of the word 'possibly' in its answer.


<sup>32</sup> Chaudri, I. (2023, April). The disappearing computer — and a world where you can take AI everywhere [Video]. TED. [https://www.ted.com/talks/imran\\_chaudhri\\_the\\_disappearing\\_computer\\_and\\_a\\_world\\_where\\_you\\_can\\_take\\_ai\\_everywhere](https://www.ted.com/talks/imran_chaudhri_the_disappearing_computer_and_a_world_where_you_can_take_ai_everywhere)

<sup>33</sup> Chaudri, I. [@imranchaudri]. (2023, 30 November). *Playing around* [Post]. X. <https://twitter.com/imranchaudri/status/1730093967969259617>



From left to right: PLAUD NOTE, Rewind Pendant, AI Pin, Tab, rabbit r1, Brilliant Labs' Frame.

There is a lot of speculation about the future of this Ai Pin online. Incidentally, the dismissal of ten of Humane's two hundred and fifty employees in January 2024 is seen as an indicator for its premature demise.<sup>34</sup> Jokingly, comparisons to the Juicero, a machine to make a glass of fruit juice, were being made. This 'Nespresso of fruit juices' as a start-up received 120-million-dollar funding from Google. As a comparison, the last commercially successful moon landing cost just a little less. The 400-dollar appliance came to the market combined with a subscription for specially designed 'produce packs' you could hang into it. Place your glass under it, switch it on, done! But it soon became clear<sup>35</sup> that you could squeeze those packs by hand over a glass and you didn't need the apparatus at all. And this is the most important question of all:




## Do we need all those new devices such as the Ai Pin anyway?



<sup>34</sup> Heath, A. (2024, 10 January). Humane lays off 4 percent of employees before releasing its Ai Pin. *The Verge*. <https://www.theverge.com/2024/1/9/24032274/humane-layoffs-ai-pin>

<sup>35</sup> Multiple videos about it can be found on YouTube, see: [https://www.google.com/search?client=firefox-b-d&sca\\_esv=43bda70055d-f04a0&hl=nl&q=Juicero&tbm=vid&source=lnms&sa=X&ved=2ahU-KEwiLo76Cj7yEAXUKhP0HHYS2AUAQ0pQJegQIDRAB&biw=1792&bih=1035&dpr=2#fpstate=ive&vld=cid:2b091aa9,vid:5lutHF5HhVA,st:0](https://www.google.com/search?client=firefox-b-d&sca_esv=43bda70055d-f04a0&hl=nl&q=Juicero&tbm=vid&source=lnms&sa=X&ved=2ahU-KEwiLo76Cj7yEAXUKhP0HHYS2AUAQ0pQJegQIDRAB&biw=1792&bih=1035&dpr=2#fpstate=ive&vld=cid:2b091aa9,vid:5lutHF5HhVA,st:0)



## The Pendant, the Tab, and PLAUD NOTE

Let's review a number of these new devices before we come to a verdict around the necessity of them. There is the Pendant of RewindAI. This intelligent necklace continuously listens to your environment and captures every conversation it intercepts, encrypted onto your phone.<sup>36</sup> Unlike the Ai Pin, the Pendant still needs your mobile phone to operate. Venture capitalist Marc Andreessen sees in RewindAI a perfect example of how AI enriches human intelligence. Sam Altman shares his enthusiasm and sees it as a strong example of personalised AI. This is preaching to the choir, as both Altman and Andreessen have invested in the company. For those concerned, Rewind states that they have a 'privacy-first' approach, and nobody's voice can be captured without their consent. How they are going to tackle this, isn't entirely clear yet.

More of these types of recording equipment will be launched, like the Tab by twenty-year-old student Avi Schiffmann.<sup>37</sup> During a presentation in a lecture room, he is wearing his Tab around his neck, and has it listen in to his speech. On the screen behind him, chat messages sent by Tab appear, giving advice on how to approach his presentation. In this way, you'll always have a real-time AI coach on hand.

Other parties are thinking about a new type of connectable too. The PLAUD NOTE comes to mind, a 'ChatGPT Empowered AI Voice Recorder' with OpenAI's Whisper technology built in. The NOTE records all your conversations, translates if needed, summarises, remembers all, and makes everything that happened on the day searchable. The device is for sale for 159 dollars, which includes a three-month PLAUD membership. You attach the NOTE to your phone for it to also record and summarise your phone conversations or slide it into your jacket pocket.

---

<sup>36</sup> <https://www.rewind.ai/pendant>

<sup>37</sup> Schiffmann, A. [@AviSchiffmann]. *I just built the world's most personal wearable AI! You can talk to Tab about anything in your life. Our [Post].* X. <https://twitter.com/AviSchiffmann/status/1708439854005321954>




## Down the Rabbit Hole with rabbit r1

Start-up rabbit markets the r1, a striking reddish-orange, squarish device for the price of 199 dollars without further membership costs. The little device is equipped with a 2.88-inch touch screen on the left and a physical scroll wheel on the right. Above this scroll wheel sits a 360-degree rotatable camera named rabbit eye. It looks a bit like a childish game console, but the intentions behind this gadget aren't childish in the least. This wannahave is lightning fast in answering questions, it completes tasks through a mysterious twist from a Large Language Model to a so-called Large Action Model, and via speech commands connects to all apps you commonly use through their 'rabbit hole' application, so different tasks are performed in conjunction.

"We do not intend to replace the smartphone," Jesse Lyu, founder and CEO of rabbit, explains during a press briefing. "A phone is primarily an entertainment device and not always the most efficient to get something done. It is a device that kills time but doesn't save time. To arrange a dinner with a co-worker, sometimes four to five different apps are needed. Large Language Models offer a universal solution for natural language, and we strive for a universal solution for services – they merely need to understand you." Jesse's ambition is to build the simplest conceivable computer, one so straightforward that you don't have to learn how to use it. He would rather not use the term 'computer' for the device: rather, it's a companion. He continues:

"And with the newest AI models (LLMs) you suddenly notice that your questions and commands are being understood."





And like that, he dismisses Alexa, Siri, and Cortana, and presents his peculiar looking rabbit r1 as their ultimate replacement.

Interaction with the r1 primarily runs via the 'push-to-talk' button that activates the rabbit OS. As soon as a user asks a question or gives a command, a stylised rabbit head appears on the screen and the r1 gets to work at lightning speed. Whether it's booking an Uber, finding a recipe for leftovers in the fridge, or identifying artists that have sampled a song by The Isley Brothers, the r1 seems to handle these tasks effortlessly, at least, according to the pre-recorded demo video. It is claimed that the r1 answers questions ten times as fast as Chat GPT.

As mentioned before, the Large Action Model of rabbit OS is at the heart of the r1, with the goal to automate simple tasks. These LAMs learn by demonstration: they observe through the camera how a person performs a task in a mobile, desktop, or cloud interface, and subsequently mimic the task. In this way, you 'train' your own army of 'rabbits' to automate specific tasks, as it were. As soon as these 'rabbits' have mastered the task, you can ask the r1 to carry it out fully automatically. Rabbit's approach closely resembles the recently emerged 'agent' type AIs. Those are machine-learning models trained on common user interfaces such as websites and apps.

## Brilliant Labs' Frame


And then there are the smart glasses ... they still have to contend with the 'glassholes' stigma that roamed around the introduction of Google Glass in 2013. We won't discuss all new models of smart glasses but mention a few: Focals by North, Ray-Ban by Meta, the smart glasses by Bose, and of SnapChat (with the latter two having been taken off the market), and the Apple Vision Pro. Adding a weight of 600 grams to your head, the last can be seen as an outlier.

## Brilliant Labs' Frame gives us some hope that smart glasses could really lead somewhere.

The operating system of the Frame, its marketing showpiece is called NOA and – you will have guessed – runs on OpenAI (GPT-4, Stability AI, and Whisper AI). In this case, the LLMs are right on your nose. The Frame offers a functionality comparable to the gadgets mentioned. As you give assignments, just like you do with Chat GPT, NOA will get to know you a bit and, just like the r1, hopes to become your companion.

## The Robot Demo War Has Broken Out

Within the outliers of AI gadgets, exciting developments are taking place. We are talking about LLM robots, or 'humanoids', machines with a human appearance that talk, listen, and observe like a human. These robots don't have to be programmed, thanks to the multimodal properties of the LLMs they perceive the world like us humans do. OpenAI



as well as NVIDIA and UBTech-Baidu have shown demos in March 2024 of what they are working on. And this has understandably been praying on all of internet's robot fans' minds since then.

In the OpenAI demo we see a man at a worktop initiating conversation with a robot. Asked the question if it would like to offer him something to eat, from all objects that are present, the robot selects an apple and gives it to him. Then the man adds some clutter to the worktop and places a basket on it as well. Subsequently, the man asks the robot to explain why it 'did what it did' and to clear up at the same time. The robot answers that the apple is the only edible thing that was in front of it and that was the reason it gave the man the apple. The robot says this while it picks up the basket and throws the clutter in it. Anyway, the demo continues for a while, and all is going well. We watched a multitasking robot that can explain its own behaviour, pick up objects, can observe, and can hold a conversation in crystal clear language.

A few days later, we see another demo by UBTech in cooperation with Baidu. It is a Chinese demo, in which the robot irons clothes and handily folds them, to then explain what he just did. And we haven't even mentioned the robots that shared the stage with NVIDIA's CEP at a recent conference. At the end of his presentation, he is even joined by two lookalike Wall-E robots that seem to have just walked out of the Disney film.

NVIDIA robots are the result of the project GROOT: a 'general purpose foundation model for robots' that is based on a 'system on a chip' (SoC). Exactly the foundation models Andrew Ng spoke about in our previous chapter. In those foundation models, NVIDIA sees a breakthrough in how we can safely and effectively cooperate with robots. The purpose: increase predictability in the behaviour of robots. A blunder in a prompt to produce an image is very different to a blunder in a prompt to a robot that is loading the dishwasher. It is clear that a lot more is involved when similar robots, provided with a quality warrant, are going to populate our living rooms.



The LLM Robot War; from left to right: OpenAI, NVIDIA, and UBTech-Baidu.



Possibly the opposite of advanced AI gadgets and humanoid robots; AI toothbrushes. They are the precursor of a tsunami of 'smart' consumer products.

These demonstrations each are undeniably impressive. The newest technological breakthroughs make them easier to get along with. The intelligent interface we see here in practice, almost makes us long to acquire a robot like it. Because this all looks very convenient. But they aren't for sale yet and, as said, we don't know yet how they will behave when they leave the laboratory and make an appearance in our own kitchens. And how will we respond, or our pets and children? For now, it is merely a demonstrative arm wrestle between the various AI giants.


## Do we need them?

On top of these advanced AI interfaces and humanoid robots, we can expect a whole range of existing products with embedded AI. Think about an artificially intelligent help while brushing your teeth, doing the laundry, or watering the plants.

So do we need them, those smart devices? If it only were as simple as a yes or a no. We'll give you a list of aspects that determine success or failure. For example, there are uses that are 'unintended' by the producer. Buyers deploy devices in a different way than the makers had intended or devised. This allows for the opening of new markets that the makers had not foreseen. Simultaneously, adoption rates can disappoint, as what the makers were

hoping for as being the most in-demand features of a product, isn't shared by consumers in practice. Then there are costs. Steve Ballmer didn't believe in the iPhone because it would be too expensive, but he was proven wrong. Sometimes, it is still a guess what price people are willing to pay for a product. Social acceptance is at least as important. The first people with mobile phones were ridiculed, the first wearers of smart glasses were scorned. The latter party because people felt their privacy being invaded as the glasses contain cameras. That danger is lurking again, as many of the new devices are quietly listening in the background and often have a camera.

Moreover, the technology needs to work as promised. TEDx talks and keynotes are nice, but not proof of widespread adoption. And finally, there is gadget blindness. The first buyers are innovators that love new technological gadgets. But a much larger group of people, the mainstream market, basically does not like technology and will only buy if there is a good reason – read: a clear application – for it. The danger of gadget blindness is that we assume a product will be successful because the 'gadget grabbers' are enthusiastic about it. We are blind to the arguments of the rest of the potential buyers, that are willing to buy a product for completely different reasons.



## Chapter 6: Conclusion and Speculative Outlook

Although the AI billions are everywhere, the AI gadget industry is running at full speed, and the agent technology seems promising to ensure the quality of new AI models, we can only speculate how the future of the intelligent interface will develop. We must keep in mind that with earlier interface revolutions, the future took surprising turns as well, and many of our former predictions turned out to be wrong. We realise that the introduction of a new interface through which people consume information, always triggers a revolution. The difficulty lies in the fact that it is nearly impossible to interpret the essence of this revolution while it is ongoing. It can only be determined afterwards. That makes it all the more complicated to give a definitive answer to important questions yet.

The question ‘Do we need them?’, which we also asked in previous chapters, does not only apply to

the new manifestations, but to a certain extent to the phenomenon of generative AI as a whole. We dwell for a moment on the faux pas by Sam Altman himself, who made it clear at Davos that the energy needs of this new AI wave cannot be satisfied with current facilities. Add the fact that we are also dealing with other forms of scarcity – we cited our own report on this several times – letting us realise that the speed of AI developments lies within the limits of physical possibilities. For anyone with big plans around AI adoption, it is good to realise that societal resistance won’t just come from the angle of privacy, intellectual property, alleged job losses, etcetera, but also from the angle of the environmental lobby, or local authorities who don’t want the AI hyperscalers in their backyards. The prediction that the ICT sector will be responsible for 14% of all carbon emissions in 2040, was made even before the LLM revolution erupted. If the AI interface would purely be devoted to ‘social good’ and solving the world’s problems, we could have outlined a merely bright future, but in reality, this intelligent interface will most likely be driven by very different motives.

This critical note puts the speculative outlook in a different light. But if we leave these issues aside for a moment, then we can delightfully daydream at the thought that dealing with information will become so much easier, that we will save time, finally get empathetic systems rather than the ‘computer-says-no’ syndrome, and that our love for technology may rise to unprecedented heights thanks to the deeper feelings we will develop for this new interface.

### The Ease of the Natural Interface

Signs suggest that our dialogue with machines will be as smooth and natural as conversations between humans. These advances are leading to a transition




from traditional graphical user interfaces (GUIs) to more intuitive and natural user interfaces (NUIs). The concept of 'search' flows into 'conversation'. The goal is to create an interaction that will feel just as natural and intuitive as speaking with another human being. As a result, the device shifts more and more into the background, becoming invisible in a way, and at the same time omnipresent.

The success of these conversations of course depends on the quality of the underlying knowledge and interpretation systems. Anyone who has ever invoked Siri or Alexa and found themselves in a Babylonian confusion of speech will know what we mean. But by the time we have all that sorted out ...

## A Companion that Wins Time Rather than Kills Time

Rather than directly communicating with websites and apps, an AI companion will take on that role, making the interaction with systems more efficient as well as more effective. The assistant will filter out any distracting content and will lead directly to the desired information or action. This results in less time spent perusing irrelevant content and more time for what really matters. A personal AI assistant then becomes the main act between the individual and the digital world. This development signals the end of the era of traditional search engines such as Google and Bing and their underlying advertising model. With the rise of generative AI, the frustrating web search can be skipped, and we receive instant, ready-made answers faster. Whether this marks the end of the search era for good, we will have to find out. Technology, like media, has a knack of lasting longer than we might think. The library, the book, the record player, and the cinema were all written off years ago, but nonetheless are still alive and kicking.



More often, we see that new inventions nestle alongside existing ones, creating a new balance. But more important than the question whether the search engine will survive, is the realisation that not just the interaction with the internet but also our everyday reality is changing, with AI becoming imperceptibly intertwined with our daily lives.

What some call a companion, others call a copilot. In a blog post, Microsoft CEO Satya Nadella talks about the 'era of copilots'. In this copilot model, man and machine work together, complementing each other's unique strengths. People will focus more on areas in which human traits such as creativity, empathy, and ethical reasoning are crucial. This is in contrast with the concept of an 'autopilot', which suggest that machines operate completely independently without human input.



## An Intermediary with Intelligence *and* Empathy

One last important observation. The way people currently use the internet to find information, is largely impersonal. It is a transactional relationship in which a browser functions as an intermediary between the persons and the task they like to complete or the content they want to consume. This intermediary is about to become more intelligent, better informed, and more empathetic. Generative AI is able to turn people's interaction with the internet from transactional to personal, making them feel more digitally understood and relevant. Along with the ability to use natural language, to get an actual answer rather than an index of links, and to ask questions, our relationship with technology is changing. Snapchat's chatbot MyAI is an early example of this trend, intended as an always-avail-

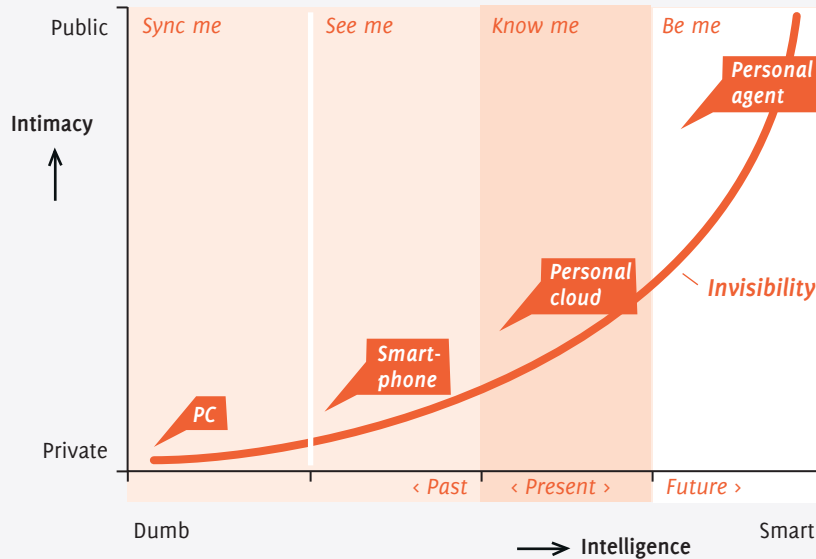
able friend that assists young people with their daily tasks and conversations. This concept is further explored by Character.AI, where users can personalise AI to take on any imaginable role, from historical figures to deceased loved ones and fictional characters, further blurring the line between reality and fantasy.

George Colony, CEO of analyst firm Forrester, describes this new world in his blog post 'The Generative AI Imperative' as follows: "You'll type in [www.taylorswift.com](http://www.taylorswift.com) and go to the site. That won't change. But what you'll find there will be very different from the old web. Yes, there will be some artistic pictures of the latest tour. But the primary experience will happen through the "Talk to Taylor" prompt box, where fans can buy tickets, buy merchandise, find out about tour dates, and get the latest videos, all through a conversation. And fans will talk to the "Taylor Language Model." These personal conversations with "Taylor" will be driven by a model trained on the artist's lyrics, interviews, statements, diaries, and proprietary content."

## Deeper Feelings for the New Interface

Replika takes this interaction one step further by allowing users to build deep, emotional – and sometimes romantic or sexual – relationships with their AI-created partners. At the other end of the spectrum are virtual influencers and streamers that draw the attention of millions of viewers despite their non-existent physical form. For example, virtual influencer Emily Pellegrini managed to turn a German professional football player's head so much that he asked her out on a date. Caryn Marjorie, a well-known social media influencer with close to 2 million followers on Snapchat, has already taken a stab at such a future and created a digital version of herself in collaboration with the company Forever Voices. This digital clone was developed by analysing





2,000 hours of Marjorie’s YouTube videos, combined with GPT-4 technology. Fans now have the option to chat with this digital version of Caryn for 1 dollar a minute. Marjorie’s goal with this project is to reduce the loneliness among her followers. It is certainly very lucrative for her. The real-life version expects to earn around 5 million dollars per month with her digital likeness.

So these GPTs will help us with a wide range of activities. Not only will they perform tasks across different applications, but they can also learn from our activities from both our professional and personal lives, making them better at recognising our preferences, intentions, and habits, and even be able to make suggestions before we ask for them. This humanisation of technology takes place in four stages.

The first three (Sync Me, See Me, and Know Me) have already been initiated. The subsequent phase (Be Me) is taking place now and in the near future. The end result will feel like the whole world is changing into a personal operating system.



## Conclusion

This speculative outlook raises fundamental questions about the nature of humanity, intelligence, ecosystems, and our dealings with, and dependence on, our environment. According to technology philosophers, technology mediates our experience of reality. And now that that technology is becoming increasingly intelligent, what does that do to our knowledge about ourselves, each other, and the world? To how we experience our surroundings? Can we see and understand more? Will it allow us to do, feel, and think fundamentally different things? And might we cooperate in different ways?

These are all questions that lay a good foundation for the next report in this 'Scarcity in Abundance' series. The fourth report, now in the making, has the draft title 'In Search of Regenerative Intelligence' and explores all kinds of intelligence, from autonomously acting AI agents in Self-directed laboratories, to human intelligence and stupidity. We will look even further and are guided by scientific insights relating to consciousness and intelligence in animals. And reflect on emergent intelligence in natural eco-systems. Can we call this natural intelligence? We will look for examples in which these forms of intelligence come together to create more value in synergy than if they function separately. Think of new forms of business with not only AI in the boardroom, but also the voice of nature. Two forms of intelligence, artificial and natural, alongside human intelligence. The aim is to outline a new perspective where organisations will function radically different and can thrive in the rapidly changing world.

## About the authors



### Menno van Doorn

Menno van Doorn is the director of SogetiLabs' Research Institute, boasting a career spanning a quarter century in future studies. His accolades include being christened the 'IT researcher of the year' by *Computable*, a distinguished IT magazine. Van Doorn's academic pursuits are deeply rooted in behavioral economics and the science of advertising.



### Sander Duivestein

Sander Duivestein is a keynote speaker, trend analyst, internet entrepreneur and strategy consultant on the impact of digital technology on people, companies, and our society. He is also a frequent guest on various radio and television programmes.




### Thijs Pepping

Thijs Pepping is a techno-philosopher focused on the relationship among people, the planet, and technology. His special interest lies in the art of living and the experiences of nature, time, and reality.



### Mike Buob

Mike Buob is the Vice President of Experience & Innovation at Sogeti. Mike has a diverse background in technology, innovation and strategy which has allowed him to play a critical role helping organisations with their transformation and innovation initiatives.



### About VINT labs.sogeti.com

VINT, the Sogeti research institute and part of SogetiLabs, provides a meaningful interpretation of the connection between business processes and new developments. In every VINT publication, a balance is struck between factual description and the intended utilisation. VINT uses this approach to inspire organisations to consider and use new technology. VINT research is done under the auspices of the Commission of Recommendation, consisting of

- K. Smaling, Chief Technology Officer Continental Europe Aegon (chairman)
- Jørgen Behrens, Vice President and General Manager Google Maps Automotive
- M. Boreel, Chief Technology Officer Sogeti Group
- Paul Dirix, Chief Executive Officer Port of Moerdijk
- L. Holierhoek, Director Offshore Energy at Royal IHC
- D. Kamst, Founder and Chief Executive Officer Klooker and Smyle
- M. Krom, Associate Anderson MacGyver
- T. van der Linden, Group Information Officer Achmea
- Prof. dr. ir. R. Maes, Professor of Information & Communication Management Academy for I & M
- P. Morley, Lecturer Computer Science University of Applied Science Leiden
- J.W.H. Ramaekers, Head of Sogeti Netherlands
- E. Schuchmann, Chief Information Officer Amsterdam UMC
- R. Visser, CIO NN Group
- J.P.E. van Waayenburg, former Head of Capgemini Netherlands

### About SogetiLabs

SogetiLabs is a community of over 150 technology leaders from Sogeti worldwide. SogetiLabs covers a wide range of digital technology expertise: from embedded software, cybersecurity, deep learning, simulation, and cloud to business information management, IoT, mobile apps, analytics, testing, and blockchain technologies. Visit [labs.sogeti.com](https://labs.sogeti.com)

### About Sogeti

Part of the Capgemini Group, Sogeti makes business value through technology for organisations that need to implement innovation at speed and want a local partner with global scale. With a hands-on culture and close proximity to its clients, Sogeti implements solutions that will help organisations work faster, better, and smarter. By combining its agility and speed of implementation through a DevOps approach, Sogeti delivers innovative solutions in quality engineering, cloud and application development, all driven by AI, data and automation. Visit us [www.sogeti.com](https://www.sogeti.com)