ACCEPTABLE AI

pelckmans.be
Brasschaatsteenweg 308, 2920 Kalmthout, Belgium

Cover design: Studio Vonk & De Boer
Interior design: Mulder van Meurs, Amsterdam
Author photo: Fotografie Vero
Translation: Heather Sills

D/2025/0055/333
ISBN 978 94 6530 050 4
NUR 800 – Business General
THEME UYQF, MBP, KJMN

Also available as:
E-book: ISBN 978 94 6530 321 5

pelckmans.be

facebook.com/pelckmans.be
instagram.com/pelckmans.be

Tom Braekeleirs

ACCEPTABLE AI

The 7 dimensions of a human-centred AI evolution

PELCKMANS

Dedicated to everyone who works in healthcare,
whatever their role.

CONTENTS

FOREWORD

'You are about to enter the greatest golden age of human possibility...'
– Sam Altman

We are at a pivotal moment. Not just in the history of technology, but in the history of humanity. And that's no exaggeration. In 2024, two Nobel Prizes were awarded to artificial intelligence applications. That event, in both what it represents and what it accomplishes, makes it clear that AI has long since outgrown the stage of hype and experimentation.

One of the recipients of the Nobel Prize in Physics was Geoffrey Hinton, pioneer of neural networks. Without his fundamental work in deep learning, today there would be no GPT, no AlphaGo and no generative AI. But what made his moment even more memorable was the warning he gave in his acceptance speech. Hinton called for caution, for a broad public debate, and for awareness of the risks that come with this powerful technology.

One of those to receive the Nobel Prize in Chemistry was Demis Hassabis, founder of DeepMind. His work on AlphaFold – which can predict the complex 3D structures of proteins with unprecedented accuracy – has opened the door to a revolutionary way of developing medicines. So revolutionary, in fact, that DeepMind's biotech division was spun off by Google into a new company: Isomorphic Labs. No longer building algorithms to win games like chess, but to better understand – and heal – life itself.

It's a sign of the times: artificial intelligence is moving from the periphery to the heart of science, medicine, economics and society. But with that shift comes another, more difficult question: what does making AI 'acceptable' actually mean?

This is not a purely technical question. It is not just about performance or scalability. It is about humanity, about trust, about responsibility. Can we build systems that are not only accurate, but also explainable? Can we integrate technology into our daily lives without compromising ourselves – or our ethical standards? How do we deal with decisions made based on models that we do not fully understand? What do we do with algorithms that learn, but can also deviate?

This book does not provide a simple yes-or-no answer. It isn't all about singing AI's praises or predicting doomsday scenarios. The author explores the many layers of what it means to develop and apply AI in the real world. He considers usage, desirability, ethics, legal boundaries, scalability and accountability. Tom dares to ask whether we really understand AI, and whether that is a requirement to be able to use it responsibly.

Underlying this is a deeper question for society: how do we deal with a technology that is no longer fully deterministic? Systems that do not always give the same response to the same input, that get creative, that sometimes even hallucinate. Is this wrong – or is it, as Hassabis suggested, perhaps a new form of originality?

What if we were to consider AI not as a tool, but as an unpredictable yet valuable partner in the way we think? Something that complements rather than replaces us. Something that increases our possibilities, but also our responsibilities.

This book explores the contours of this new relationship between humanity and technology. Not with the intention of providing conclusive answers, but of asking the right questions.

The question at the heart of this book is surprisingly simple: will AI become a normal technology?

Some think so. The influential paper 'AI as Normal Technology'[1] argues that we should not overestimate AI. The revolution is a gradual one. Just as electricity once changed everything but has since become so commonplace that we now take it for granted, AI will also eventually become normal.

But one counterargument is AI 2027, a scenario analysis stating that we will be confronted with a fundamental reshaping of the world, driven by AI agents that can plan, reason and negotiate independently. This is not science fiction. It is a logical consequence of the exponential growth we are experiencing today.

This tension – between sober realism and awe of the unknown – will continue to generate much debate in the coming years.

What we do know is that we have a responsibility to shape the future of this technology. To not leave the debate to tech companies alone, but to claim social ownership of it.

If we can do that – with vision, vigilance and imagination – then we can also realise the promise of AI. Perhaps we are indeed on the threshold of the 'greatest golden age of human possibility'. It is up to us to make it acceptable. And valuable.

Peter Hinssen
June 2025

INTRODUCTION

It was a dry and somewhat windy Wednesday in San Francisco, with November drawing to a close and the autumn sun doing its best to banish the greyness. Inside the city's many office buildings, something was brewing. One small office on 18th Street was a hive of activity. Looking at the grey facade of the building at number 3180, you would never have guessed that inside, a team of developers were putting the finishing touches on the next revolution. This is something you often hear in the start-up hub of this Californian city. Unaware of the storm that their work would unleash, the team pressed on, fuelled by coffee and foosball tables.

A team of engineers sat at long wooden tables, programming on laptops with too many stickers; half the screen covered in code, the other half in a Slack conversation – or forty. At 6 pm, one of them hit the enter key. And that was it. No Tim Cook speech with epic intro music. No clumsy Zuckerberg in an oversized VR headset. Just some lukewarm pizza and a modest ripple of applause filled the building. It was just… there. Online. As if someone had nudged open the door to the future with the words: 'Come and take a look.' The screen blinking in front of co-founder and CEO Sam Altman showed just four words: 'ChatGPT is now live.' One simple line. No exclamation marks. Yet the world would never be the same again.

OpenAI was founded in 2015 as a research group with the goal of making AI accessible to the whole world. Seven years later, they launched a new model. It wasn't a big release, just

the next version of the model, known as a point release: 3.5. But somewhere along the way they had the brilliant idea to let people converse with the model. There was nothing new about chatting; from ICQ to MSN Messenger, internet history is peppered with chat programs. But until that point, there was always a human on the other end of the line. This time, things would be different. The virtual conversation partner was born. And it conquered the world at breakneck speed. Within hours of its launch, the first tweets appeared from people who had tried it out. The news spread faster than a cold in kindergarten. Journalists, nerds, students, and everyone in between started using it. Some were seriously impressed, others seriously concerned, but nobody was immune to it. What was clear was that we had entered a new era. The wonderful world of generative AI was becoming accessible to everyone. At an unprecedented pace, it was taking its place in the universe of the internet. The last time a technology had claimed such a revolutionary spot was the arrival of the internet itself. But to experience it, you needed a connection to… the internet, which everyone now had. In a matter of weeks, they reached the milestone of more than 100 million users, pulverising the record of both TikTok (9 months) and Instagram (2.5 years). A new undeniable force had claimed its seat at the technology table. Putting the genie back in the bottle was no longer an option.

The story of OpenAI and their ChatGPT brought focus to the discussion around AI. Suddenly it hit the global news and became a topic in many boardrooms – a completely new item on the agenda for many organisations. Before that, it had been seen as a technical gadget, a tool for the IT department…

However, AI is not new. The term is actually more than seventy years old. The basic concepts of modern AI systems are considered old, with the foundations of the great revolution

we refer to as generative AI having already been laid in the 1980s. Yet it has still taken us by storm. The rapid progress with which AI has entered our lives suddenly became clear: a dream come true for some, a nightmare for others. The fact is that it has become a reality that we have to learn to deal with. But how? How can we, as an organisation, best use AI? It is impossible to keep up with the number of courses and training programmes on how to use (Gen)AI. Even educational institutions are asking themselves how they should move forward. Does writing an essay still make sense now that every student has access to a tool that can spit out grammatically correct texts at an unprecedented speed, and (sometimes) of unprecedented quality?

I work in healthcare. This has been my world for a decade now and the impact of AI is also palpable there. But far too often I notice that people – and not just in healthcare – have a kind of conservative reflex, more inspired by a lack of knowledge and fear of the unknown, than a fear that machines will suddenly take over the world. I refuse to go along with the argument that 'the rise of the machines' is underway and that they're coming to wipe out humanity. There is currently no evidence for this, and it largely arises from our own human idea of power and control, which we also want to project onto machines. It's already anthropomorphic enough.

At the same time, we are often on a roller coaster of hype and overenthusiasm. The launch of every new technology is accompanied by bold statements about how this or that will completely disrupt sector X, Y or Z and finally solve the world's major issues. Every new model is groundbreaking. Suddenly everyone is an expert, and the unbridled enthusiasm and expectations significantly exceed the feasibility of the technology itself.

But that was also the reason behind writing this book. I am entirely convinced that a lot of the fear can be dispelled by knowledge, by better understanding what AI is, and also what it is not. With a better understanding of this, we can also more accurately estimate the impact on our sector, be it healthcare or otherwise. Knowledge is power. Understanding that AI is all about statistics, and is therefore driven by probabilities and complex mathematics, can very often explain why models make mistakes. It can also reassure us that complete certainty is impossible. It's about statistics.

At the same time, if we understand what AI is, what it can do and how we can use it in a responsible way, the enthusiasm might be tempered somewhat. Yes, AI is a potentially disruptive technology, but it also requires the right implementation frameworks.

I hope this book can contribute to the debate. It is by no means an attempt to preach one single truth about AI. This book is the result of my own search for the impact of AI technology and the evolution it has undergone from the first attempts to date. The danger of writing a book about AI is that it is quickly out of date. This is also why I prefer to focus on the impact this technology can have on people, and less so on how it technically works. I focus more on the 'human stack' than the tech stack. But doing the first justice does require some understanding of the latter. The basic question I started from is: what needs to happen before people can start to trust AI? What are the basic questions we need to ask in order to feel reassured? Reassured that AI does not want to take over the world, but, above all, reassured that it'll be okay – provided that a number of rules are respected. As humanity, we now have the choice to write the rules of the game, together with AI. It would be naive to think that

we can do this all by ourselves. We have to keep talking to each other, and now to AI as well. Because this technology is starting to confront us with the limits of our own abilities and the errors in our own ways of thinking. Together, we can figure it out. We can embrace AI. We can accept AI, and in doing so, move towards acceptable AI.

Tom Braekeleirs
June 2025

1

AND THEN THERE WAS AI

THE WORLD WE LIVE IN

There's no more denying it – these are interesting times we're living in. And some see this more clearly than others. One day, when we're old and grey, we'll look back on this period of transformation and realise that a lot has changed in a very short space of time. It was supposedly Lenin who said: 'There are decades where nothing happens; and there are weeks where decades happen.' We have all experienced this kind of pivotal moment, very consciously watching it unfold. The introduction of ChatGPT has been a massive wake-up call, in a positive sense, but it has also made us aware of the fact that we still have a lot of questions about how best to deal with these kinds of technologies. And we do not yet have the answers. All we can do is try our hardest to fit it into the Zeitgeist as best as possible, with a (hopefully solid) framework around it. Because the times we're living in don't want to wait. This is not an era that looks patiently at humanity and asks for permission. In the wise words of Satya Nadella, CEO of Microsoft: 'Our industry does not respect tradition, it only respects innovation.'

These are digital times we're living in. Anyone with teenagers at home knows this all too well. Personally, I've lost count of the number of discussions about whether or not there's too much screen time, whether they're continually staring at their smartphones, how bad it is for their eyes, and whether they should be playing outside instead… It seems as if our parenting skills are proportionally linked to the strength of our Wi-Fi signal. You could say we're living in an era that is perhaps best described as 'hypertime'. Everything seems to be speeding

up. Whether that is true or not is a question I'll refrain from answering for a moment. If you were born in 1900, then you lived through the arrival of aviation, space travel, the invention of television – both black and white and colour, the rise of radio, two world wars, the car, perhaps even the internet and much more. So are things really going that fast now? I think so. And especially because our brains are now also experiencing something of an overload. We can no longer keep up with the cognitive load that a multitude of impulses is asking us to process. Combine that with the amount of information produced every day and the figures are staggering. We're talking about petabytes of information being generated by humanity as a whole. One petabyte is 10^{15} (ten to the power of fifteen) bytes of digital information, or 1 million gigabytes, or for the old school among us: 745 million floppy disks. In the coming years we will increasingly start talking about exa-, zetta- and yottabytes. As a result, it seems that we have never had so little time, never had so much information available to us that we cannot process, and that our eternal thirst for knowledge is far from being satisfied. The more we know what we currently know, the more we realise that it will be a long time before we know everything there is to know. I'm well aware that I used 'know' an awful lot in that sentence, and I do not want this to turn into a philosophical book, but we do have to be aware of the context we are living in.

That context is not only digital, but also extremely turbulent. It feels as if in recent years we've been lurching from one crisis to the next without hitting the pause button. There are essentially two major types of crisis. The first are the seismic shocks. Things that happen very suddenly, that have a major impact and that also affect us personally. Covid-19 was certainly one of these. The term was also particularly apt during the eruption of Eyjafjallajökull, an Icelandic volcano that spewed out

a gigantic ash cloud, severely disrupting transatlantic air traffic for several weeks in 2010. The economic damage ran into the billions, 100,000 flights were cancelled, and millions of travellers were left stranded on the wrong side of the ocean. A volcanic eruption is not always spectacular, but as soon as it impacted us personally, it became, quite literally, a seismic shock. Then there are the undercurrents. These are things that ripple under the surface, that emerge slowly and that don't seem to be evolving all that dramatically if you look at each step individually, but that could one day have a major impact. Such as global warming to name one. This is happening very gradually. You won't suddenly wake up one morning to find that the average temperature has climbed a few degrees and that the sea level has risen by a metre. However, the consequences are becoming clearer every year. In 2023, heat waves in Europe claimed the lives of almost 50,000 people.[2] And the prognosis isn't good. It is creeping up a little at a time, a bit more every year. This kind of change takes place extremely slowly, until it doesn't. Because when it does change suddenly, it becomes a seismic shock. Undercurrents are shocks with patience, if you will. And unfortunately, there are quite a few of these in healthcare…