Foreword

This book is rare in several ways and unique in its combination of attributes. First, few, if any, books are aimed specifically at the CEO or senior executive who wants to learn about artificial intelligence. Fewer still are so admirably concise.

The content is unique in that it combines AI, knowledge and knowledge work, and business strategy. There are books that relate AI to strategy – I wrote one of them – but none that relate it to knowledge and knowledge work. If there were any, I would be aware of them, because I did a lot of work in knowledge management ten to twenty years ago, and I keep an eye out for anything that brings that topic into present times. This is somewhat curious, as AI is about embedding knowledge into computer systems, and often the target is the human knowledge worker, both for extracting their knowledge and for automating their key tasks.

Finally, the book is unique in that it looks simultaneously into the history and the future of AI. Unlike most books on AI, the author, Professor Dörfler, has worked in the AI field for many years and is familiar with the many winters and springs the field has endured. It is truly unusual, for example, to find an author who can compare expert systems to deep learning neural networks.

Why, you ask, should a busy senior executive be interested in learning about AI's past? Only because AI's past is likely to be its future – once again. Neural networks were given up for dead in the 1960s, then arose in the late 1980s and early 1990s for fraud identification in financial services. By the turn of the century, the world had largely lost interest in them again, though financial services firms still made very effective use of them. They arose again in the 2010s in the form of deep learning neural networks. Now these are perhaps the single most popular technique in all of AI. There were small advances that enabled their rebirth each time, but the basic principles remained the same.

The same historical rebirth could also happen with rule-based expert systems, a technology with which the author is quite familiar because he created a tool for developing them. Expert systems were the AI technology in the late 1980s and early 1990s. Like neural networks, they beat on in the hearts of insurance underwriting systems, healthcare clinical decision support, and multiple other applications. Yet to the press and many AI experts, they were declared dead long ago.

Dörfler is well aware of the strengths and weaknesses of logic-based expert systems, and he makes a good case that they will reappear in the future of AI. Some of the greatest AI researchers in the world have acknowledged that deep learning models are not enough to power the many future uses of AI, and the simplicity, causality, and explainability of rule-based systems are likely to bring about their return.

Don't read this book if you are only looking for what is most fashionable about AI. But if you desire to know the essence of AI – how it can capture, analyze, and apply knowledge to problems that humans need to solve – you have picked up the correct little tome. You will gain much wisdom about knowledge and intelligence – human and artificial – if you keep it in your hands or on your screen and turn its pages until the end.

Thomas H. Davenport
Distinguished Professor, Babson College and Visiting Professor,
Oxford Saïd Business School
Fellow, MIT Initiative on the Digital Economy
and Senior Advisor to Deloitte AI
Author of *Competing on Analytics* and *The AI Advantage*

1 Introduction: The Strategic Landscape of AI

Prediction is very difficult, especially if it's about the future!

Niels Bohr[1]

There is little doubt today about the significance of artificial intelligence (AI), but it is less clear why it should have a place in a series on business strategy. The book's title suggests that there is something that every CEO should know about AI – does this mean even those who do not intend to use AI? Definitely yes. And in this short book I explain why.

To avoid any confusion, it is important to clarify my position from the outset. I am an AI enthusiast, and I believe that AI can be an incredibly useful thing if we figure out how to use it well. However, I am not a carnival barker[2] who just wants everyone's buy-in, regardless of whether AI is useful for the particular topic or not. Conversely, I am also not a harbinger of doom foretelling how AI will overtake the world and enslave or exterminate humankind. This book depicts the limitations as well as the strengths of AI and delineates facts from beliefs in order to make it possible for CEOs to decide in particular situations whether they need AI and how to use it.

1.1 AI, Strategy, and the CEO

How do we know that AI is important? As with not only any technology but any phenomenon in the world, we can look for the money. AI seems to be the most expensive human endeavor ever; the four most prominent hubs – the US, the UK, Canada, and Israel – together spent well over a trillion dollars on AI, and China, India, and Russia are probably not far behind. Chengwei Liu suggested in *AoM Insight* (2020) that the first trillionaire in the world would come about from the AI business. We can also see that AI is ubiquitous; it permeates all areas of our work as well as our private lives. This takes us to the very reason why every CEO needs to know something about AI: It is not possible to opt out. In 1997, on the Oracle Co. website, a statement appeared: "In five years nobody will call it e-business. It will simply be business." It took a little longer, but it happened. Similarly, AI is not optional; if a company does not adopt it, its competitors will, customers may expect it, etc., so not using AI does not mean one is not affected by it.

[1] This saying is perhaps most often attributed to the Nobel Prize–winning physicist, but variants also appear with attributions to Samuel Goldwyn, K. K. Steincke, Robert Storm Petersen, Yogi Berra, Mark Twain, and even Nostradamus. It is also quoted as "an old Danish proverb," but it seems to be a proverb that cannot be traced.

[2] "a . . . person who attempts to attract patrons to entertainment events, such as a circus or funfair" (Wikipedia: barker).

In contrast to what we may learn from (technology) management textbooks or academic journal papers, in real-world organizations CEOs seem to be invested in ICT (information communication technology) decisions. So why is a CIO, even if there is one, not in charge of an ERP (enterprise resource planning) decision? The reason is that such decisions are never mere technical decisions. The CEO of British Petroleum said a couple of decades ago that in ICT investments, about 20 percent of the cost is the hardware and the software and 80 percent is learning and cultural change. That is, such decisions affect businesses at their very core; it changes how they see themselves and their environments. In other words, it affects their Theory of the Firm (ToF). Thus what each CEO should know about AI is how it affects their ToF and the actions stemming from their ToF – their business strategy.

The purpose of this book is to deliver exactly that. While CEOs do not need to understand the technical details of AI, they can benefit from understanding what sorts of AI there are and what each sort can be expected to deliver. In other words, it is useful if a CEO is able to see through the too often overly enthusiastic dreams of AI researchers, as well as the excessive claims of AI vendors, in order to make informed strategic decisions – as well as to decide in broad terms how to go about the implementation. The final decisions can most meaningfully be taken at the deep levels of the corporate hierarchy, where the relevant knowledge resides. This is in line with Charles Handy's (2015) principle of inverse delegation, meaning that the decisions belong where the knowledge is, and if help is needed, the specialists delegate to their bosses.

This book is not only aimed at CEOs; anyone working with and for a CEO can benefit from it. This means that AI researchers and AI vendors will also benefit from it, as they will learn about the context and potential values as well as the pitfalls of AI. AI vendors cannot decide for a client where and for what purpose it needs AI, as they do not understand the primary business processes of their clients in depth. Clients also have problems with these decisions, as they do not know what AI can deliver. This is the age-old problem of ICT implementation; what is sold is different from what is purchased. The vendor sells the capabilities of the product, while clients purchase solutions to their problems. For a successful implementation, these two need to be brought together, which will require a collaboration between a subject expert from the client and an AI expert from the vendor. In order to develop a conceptual and cultural framework in which this is possible, it is useful if the CEO is able to distinguish between facts and beliefs in AI, thus understanding the value AI can realistically deliver.

As indicated above, besides CEOs, this book is primarily relevant to those who work with CEOs. However, in a sense, all of us work with CEOs, and most of this book is about understanding how AI affects our lives, in a strategic sense,

and how to make the best use of it. In this sense, this book is a brief summary of what we need to think through when we have the chance to use an AI solution and how to use it for the greatest benefit – whatever that may be. Therefore, this book is not about giving the readers the answers but about showing how they can obtain their own answers. Unfortunately – or perhaps fortunately – generic, universally applicable solutions do not exist. As in most cases with reference to strategy, everyone must find their own.

1.2 Uncertainty: The Context of Strategy

To set the scene for thinking strategically about AI, we need to depict the landscape in which businesses exist as *uncertainty*, as reconceptualized by Frank Knight (1921) and as extended and enriched by J. C. Spender over the past fifty years (Spender, 2021). Uncertainty is characterized by knowledge absences (Spender, 2014), and these knowledge absences create opportunity spaces (Spender, 2021). The essence of entrepreneurial and managerial work is to construct a language that makes engaging with the opportunity spaces possible, rendering them actionable and manageable. In this context, AI is useful if it helps us wrestle with these knowledge absences.

Knight introduced the concept of uncertainty to contrast it with the concept of risk. In the case of risk, although we do not know exactly what will happen as a consequence of a decision, we have complete knowledge of all the alternatives that may happen as well as of the probabilities of each of the alternatives, which together add up to 100 percent; in other words, the probability distribution is known. In uncertainty, we do not have such complete knowledge; we have knowledge absences. These knowledge absences can be of different kinds.

Martin Shubik (1954) distinguished two aspects of uncertainty: *ignorance* and *indeterminacy*. Ignorance stands for what could in principle be known, only the CEO does not know. It could be information that exists but is not accessible to the CEO, or it could be infeasible or even impossible to obtain, as there is a great deal of information. In contrast, indeterminacy brings in the other players. Indeterminacy is not about knowledge that is difficult to access or obtain; the knowledge, in this case, does not exist. Typical examples are the competitors, who pursue their respective courses of action and whose behavior is not fully determined by the rules of the game – which leads to the area of game theory,[3] which is Shubik's primary

[3] Game theory is a discipline of applied mathematics, pioneered by John von Neumann (von Neumann & Morgenstern, 1953), that depicts strategic situations in economics, politics, etc. as games, i.e. using mathematical models that involve multiple (rational) decision makers, referred to as players, that make choices independently of each other but in response to each other, and the outcomes of their choices are affected by the other players' choices.

field.[4] Indeterminacy might more colloquially be thought of as *free will*. Spender (2014) further extended the model by adding *incommensurability*, meaning that the players, as well as knowledge items, cannot be compared, as they are so different that we even lack the basis for comparison. Incommensurability could also be termed *uniqueness*. More recently, Spender (2018) identified another aspect of uncertainty, which he calls *irrelevance*. The core argument here is that the information that we can express in a (formal) language may be less relevant to the phenomenon at hand and our decision regarding it than what we cannot verbalize. This means that the tacit dimension of personal knowledge (Polányi, 1962a, 1966b) is more relevant than explicit knowledge, which is why this aspect of uncertainty might also be called *tacit knowledge*.

In order to improve the orientation around the concepts of uncertainty, my coauthor Alina Bas and I (Dörfler & Bas, unpublished) organized these concepts into three realms: *known*, *unknown*, and *unknowable*.[5] *Known* refers to complete knowledge, therefore it includes risk as well as the special case of risk – certainty, where there is only one alternative with 100 percent probability. Importantly, risk is not less known than certainty; in both cases there is complete knowledge, but in "certainty" it is a complete knowledge of a specific outcome, while in the case of "risk" it is a complete knowledge of the probability distribution. Unknown and unknowable together comprise uncertainty, which is bound by the knowns. Spender notes that the knowns are also not as well-behaved as some may wish, as they can be "surprising, ambiguous, anomalous, inconsistent, or contradictory" (Spender, 2021: 129). *Unknown* corresponds to ignorance; that which is possible to know, but which the CEO just does not know here and now. *Unknowable*, in contrast, means that it is impossible to know; it covers uniqueness, free will, and tacit knowledge (or in Spender's terminology incommensurability, indeterminacy, and irrelevance).

Importantly, uncertainty does not mean that the CEO does not know anything. As Shubik wrote: "It is fairly obvious that there is rarely an economic situation in which the entrepreneur has no information at all. If such a situation

[4] Spender warns that "game theory only offers rigorous solutions when there is no uncertainty left, when each player's moves can be 'gamed out'" (Spender, 2021: 128). For us this is important, as we cannot simply leave this up to AI, as this can perhaps work well in a simulation, but in the context of the CEO's reality, there is always some uncertainty left, so a rigorous solution is impossible.

[5] This is not the same as Donald Rumsfeld's "known unknowns" and "unknown unknowns"; those are both unknowns. Sometimes we may know of what it is what we do not know, at other times we may not, but in both cases it is not a matter of learning it; the knowledge is unobtainable.

exists, then the first action of the entrepreneur must be to obtain information" (Shubik, 1954: 632).

For instance, the CEO may know about alternatives ***A***-***B***-***C***, and may know that ***A*** is more likely than ***B*** and ***B*** is more likely than ***C***, without being able to allocate exact percentages, and there may be further alternatives beyond the CEO's grasp and some of these may be more likely than ***A***. In most cases it is also reasonably safe to assume that the tax regulation, for instance, will not change; that is, it is known. Perhaps even more importantly, more than one of the four aspects of uncertainty can, and usually do, affect the CEO at the same time: There is often inaccessible information or information that cannot be processed sufficiently quickly (ignorance), there are often several or even many players (uniqueness) making their own decisions (free will), and the CEO has a feel regarding where the industry is going, the mood of the competitors, etc. (tacit knowledge).

Elsewhere we suggested that AI can be a good tool for the unknown, but to cope with the unknowable, CEOs use their intuition (Dörfler & Bas, 2020b). We can refine this argument. AI can be exceptionally useful in the realm of the unknown, but it cannot do the job on its own. As Spender eloquently put it: "Firms are contexts in which data-driven theorizing is subordinated to entrepreneurial judgment. . . . As we collide with the uncertainties that arrest our activity, we respond with imagination rather than with reason" (Spender, 2021: 125).

At the end of the day, the CEO will still need to make decisions, but AI can provide useful information for those decisions if we figure out how to use it in a sensible way. In the realm of the unknowable, CEOs will have to rely on their intuition, but AI may be useful, for example, in helping to identify the scope for intuiting, as long as it is a unique AI setup for the specific organization at hand and subordinated to the CEO's value judgment. I return to these points in Section 7.

1.3 Why This Book?

In 2017 the MIT Sloan Management Review and the Boston Consulting Group together surveyed 3,000 executives, managers, and analysts across a variety of industries and conducted in-depth interviews with 30+ technology experts and executives (Ransbotham et al., 2017). They found the following:

> The gap between ambition and execution is large at most companies. Three-quarters of executives believe AI will enable their companies to move into new businesses. Almost 85% believe AI will allow their companies to obtain or sustain a competitive advantage. But only about one in five companies has

> incorporated AI in some offerings or processes. Only one in 20 companies has extensively incorporated AI in offerings or processes. Less than 39% of all companies have an AI strategy in place. The largest companies – those with at least 100,000 employees – are the most likely to have an AI strategy, but only half have one. (Ransbotham et al., 2017: 1)

This issue is addressed within this book, which is somewhat but not quite entirely unique. This book is about the most advanced technology, yet it is nontechnological. It is primarily written for top executives (and my MBA class) so that they can learn about technology, but it will be equally useful for technical people, as they will understand the business strategy mindset a little bit better within the context of AI. There are numerous books and papers that explore the use of AI by business organizations, but they usually focus on what is available, and often within the context of one narrow functional area, such as finance or marketing. In contrast, this book approaches the organization as a whole and focuses on the strategic thinking process of making use of AI. The book that comes closest to this approach is Thomas Davenport's (2018) *The AI Advantage*; however, that book focuses more on the analytical aspect, and therefore the two can be regarded as complementary.

This book depicts a particular picture of AI at a relatively high level of abstraction, often exaggerating or simplifying issues in order to make them easier to understand – but carefully, so as to be correct in principle. As such, it offers more of a "big picture painted with a thick brush and broad brushstrokes"; that is, it is a comprehensive picture that lacks many details but from which the picture can be recognized. For this reason, although the book talks about artificial intelligence throughout, there is perhaps more psychology and philosophy than technology included. As often happens with relatively new and incompletely understood ideas still in development, in order to understand what they are, it is useful to delineate them from and even contrast them to what they are not. As the area is complex and changing fast, many details are perhaps omitted from the text and not all details are up to date; however, the details that become available will easily find their ways to anyone's attention, so this is not as important. The most important role of such "big-picture" books is to help the reader develop their own thinking about these issues in the way that suits them. Many arguments are offered, but more important are the ways in which they are constructed; showing this helps the reader construct their own argument. Having the big picture as a reference point, readers will find it easy to make sense not only of what is already there about AI but also what will come about in the future.

Having written this book for CEOs does not only refer to the content but also to its style. The text is organized into five substantive sections, each consisting of five subsections. Four subsections in each section are around 1,000 words, and each can mostly be read on its own. This does not mean that there is no

progression in the book or that the subsections do not connect – they do, but they can be read individually or a couple of them at a time and then be continued later. The fifth subsection in each section is called "So What?"; these are sort of takeaway points; something you can take with you if you plan to discuss AI or to serve as a brief recap. The reason for designing the text this way was to try to adapt to the busy schedules of my readers.

The first section, "Ex Machina," is the most technical one; this introduces the different types of AI. It is important for CEOs to have a nuanced and complete picture regarding what is out there. The next three sections deal with the most important human aspects of what we associate with intelligence; knowing, learning, and creating. Each compares and contrasts human ways with machine ways. The most significant takeaways from each section are the aspects on deciding what to use AI for and where we need human experts. The last section focuses on the ethical aspects of AI. This is the most dynamic area of AI at this time; we come across new moral dilemmas associated with AI on a daily basis. Therefore, by the time the reader gets into this book, there will be novel moral issues not covered – but the hope is to provide a solid basis for tackling these mindfully. I also believe it is in the area of ethics, rather than technology, where we will learn the most about AI over the next decade or so. In Section 7, I will get back to the issues that I brought up in Section 1: namely, how AI can help CEOs cope with uncertainty.

2 Ex Machina

Any sufficiently advanced technology is indistinguishable from magic.
Arthur C. Clarke[6]

This section focuses on the different types of AI; each is explained in its respective historical context. The reason for including this historical perspective is that it helps us to understand what the particular AI types were supposed to accomplish, how they were expected to work, and what was their expected scope of validity. In other words, we should understand the assumptions behind the particular AI types and see what has changed since their inception. This section also helps us to develop a working AI vocabulary, demystifying concepts like "deep learning" and explaining the technological background in fairly understandable, nontechnical terms. Although this is not necessary for CEOs in order to think about AI, it is very useful when it comes to communicating about AI, particularly if it is with people who are "high tech native," who were born familiar with today's information technology. An important complementary read to this section is Pamela McCorduck's (2004) personal review of AI history, which is rooted in interviews with many of the early AI gurus. I see things somewhat differently and make my case below, but the richness of McCorduck's story is admirable. This section also retells some history, elaborating on the origins of the different types of AI and helping to uncover the underlying assumptions about them.

2.1 Symbolic Reasoning Systems

According to AI folklore, the history of AI started in January 1956 when Herbert Simon taught his first class on Mathematical Models in the Social Sciences at the Carnegie Institute of Technology (today Carnegie-Mellon University), having just returned from the New Year holiday. It was a working holiday, as he told his students: "Over the Christmas holiday, Al Newell and I invented a thinking machine" (Simon, 1991: 206). One of the students in the class was Edward Feigenbaum, who recounts that this sentence made him so excited about the topic that he became Simon's PhD student (Feigenbaum, 1992: 3). Later we saw Feigenbaum become the father of another type of AI, but for now, we stick with Newell and Simon.

The "thinking machine" that Newell and Simon created, together with Cliff Shaw from the RAND Corporation, was called the Logic Theorist, or Logic Theory Machine (Newell & Simon, 1956). Just to hint about the state of technology at the time, Newell and Simon had to fly to Los Angeles to do this

[6] This statement became popular as Clarke's Third Law (originally published in Clarke, 1962).