




THE EXECUTIVE GUIDE TO ARTIFICIAL INTELLIGENCE

Cutting Through the Hype - How to get
the most from AI in your Organization

SECOND EDITION

ANDREW BURGESS



The Executive Guide to Artificial Intelligence

Andrew Burgess

The Executive Guide to Artificial Intelligence

Cutting Through the Hype - How to get
the most from AI in your Organization

Second Edition

palgrave
macmillan

Andrew Burgess
Greenhouse Intelligence Ltd
London, UK

ISBN 978-3-031-50721-2 ISBN 978-3-031-50722-9 (eBook)
<https://doi.org/10.1007/978-3-031-50722-9>

1st edition: © The Editor(s) (if applicable) and The Author(s) 2018

2nd edition: © The Editor(s) (if applicable) and The Author(s), under exclusive license to Springer
Nature Switzerland AG 2024

This work is subject to copyright. All rights are solely and exclusively licensed by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Cover illustration: Ukususha/iStock/Getty Images Plus

This Palgrave Macmillan imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

If disposing of this product, please recycle the paper.

This book's second edition is dedicated to my father, Ken Burgess, who passed away in 2022, and my mother, Annette Burgess. I owe them so much.

Preface

The first edition of this book was published in 2018, during what seemed like the peak of AI hype. We had seen huge advances in how AI was being used in research and development, particularly with the advent of Deep Learning, and we were seeing AI being used to transform processes and businesses at the ground level. I wrote the book specifically to help non-technical business executives make the most of these technologies without getting carried away by all that marketing hyperbole and froth. But all that hype seems like mild enthusiasm compared to the hysteria we experienced following the launch of ChatGPT. Although generative AI models had been around for a few years, ChatGPT put AI firmly into the public consciousness. Here was an advanced AI technology that anyone could use, providing information and insights beyond anything they had experienced before, all through their humble web browser. This was the ‘democratisation of AI’ in action. But it also brought risks, which some claimed were existential, which only heightened the levels of chatter in the press. So the need has arisen again—to write a book about AI that cuts through the hype, explaining to executives the things they should care about and things they can ignore. This second edition of *The Executive Guide to Artificial Intelligence* is more than simply a tweaked version of the original—many parts have had to be rewritten in full to accommodate the huge advances and opportunities that Generative AI has brought to

the world. It's an exciting time to make use of these technologies, but also a worrying one as we grapple with the ethical issues they bring. I hope this book helps you navigate these challenges so you and your business can flourish in this brave new world.

London, UK
2024

Andrew Burgess

Contents

1	Don't Believe the Hype	1
2	Why Now?	9
3	AI Capabilities Framework	23
4	Associated Technologies	47
5	AI in Action	61
6	Starting an AI Journey	75
7	AI Prototyping	97
8	What Could Possibly Go Wrong?	109
9	Industrialising AI	125
10	Living with AI	151
11	Where Next for AI?	163
	Index	173

About the Author

Andrew Burgess A management consultant, author and speaker with over 30 years' experience, Andrew is considered an authority on Artificial Intelligence. Through his business, GreenhouseAI, he advises organisations on their AI strategy, its application in business and its ethics, and is regularly invited to speak at conferences on these subjects. He is Visiting Senior Fellow in AI and RPA at Loughborough University and Expert-in-Residence for AI at Imperial College's Enterprise Lab. He is a member of the AI Committee at the British Standards Institute.

List of Figures

Fig. 2.1	Neural network layers	18
Fig. 2.2	Neural network weights	19
Fig. 2.3	Neural network inference	19
Fig. 3.1	AI objectives	24
Fig. 3.2	The AI framework	44
Fig. 4.1	Human in the loop	59
Fig. 4.2	Human in the loop active learning	60
Fig. 4.3	Tagging using crowdsourcing	60
Fig. 6.1	Developing an AI strategy	77
Fig. 6.2	AI maturity matrix	83
Fig. 6.3	Mapping AI opportunities	86
Fig. 6.4	AI heatmap	87
Fig. 9.1	AI playbook	127



1

Don't Believe the Hype

Introduction

Read any current affairs newspaper, magazine or journal, and you will likely find an article on artificial intelligence (AI), usually decrying how the 'robots are taking over' and how this mysterious technology is the biggest risk to humanity since the nuclear bomb was invented. Meanwhile, companies creating AI applications make grand claims for their technology, explaining how it will change peoples' lives whilst obfuscating any real value in a mist of marketing hyperbole. And then there is the actual technology itself—a chimaera of mathematics, data and computers—that appears to be a black art to anyone outside the world of technology. The dramatic introduction of Generative AI models, such as ChatGPT, has taken this hype to whole new levels. No wonder business executives are confused about what AI can do for their business. What exactly is AI? What does it do? How will it benefit my business? Where do I start? All of these are valid questions that this book seeks to address directly.

Artificial Intelligence, in its broadest sense, is already fundamentally impacting how we do business. Of that, there is no doubt. It is changing how we make decisions and enabling completely new business models to be created. It is allowing us to do things we never thought possible just a few years ago: Artificial Intelligence can identify cancerous tumours with greater accuracy than human radiologists; it can identify fraudulent credit card behaviour before it happens; it can drive cars without drivers; it can analyse huge amounts of documents in seconds; it can predict when customers (and employees) are going to desert you and, most importantly, it can learn and

evolve based on its own experiences. If that wasn't enough, this technological revolution accelerated beyond anyone's expectations in late 2022 with the coming-of-age of Generative AI. Suddenly, here was a technology that brought enormous capabilities and risks in equal measure. Some of the hype that ChatGPT and its ilk have created is absolutely justified—these are truly amazing pieces of technology that are changing how people do business and interact with information. We can simply ask for ideas, strategies, poetry, computer code and images on almost any subject we want by typing our request into our browser. But these seemingly magical abilities have led some to anthropomorphise the Generative AI models, ascribing them with inappropriate levels of consciousness (the appropriate conscious level of any AI model right now is zero). Many people are worried that AI will replace some or most of the work currently being done by knowledge workers. AI, and Generative AI in particular, provides opportunities for bad actors to maliciously influence how we think and to do it at an unprecedented scale. AI is both a huge opportunity and an ominous threat wrapped up in a bewildering bundle of algorithms and jargon.

A persistent challenge, therefore, for executives who want to get to grips with AI is where to find all the relevant and useful information without resorting to fanciful articles, listening to vendor hyperbole or trying to understand the algorithms themselves. AI is firmly in the arena of 'conscious unknowns'—we know that we don't know enough. Until business executives understand what AI is, in simple enough terms, and how it can help their business, it will never reach its full potential. Those with the foresight to use and exploit AI technologies are the ones who know what it can do and understand what they need to do to get things going. Providing you with these insights is the mission of this book. Over the course of the eleven chapters, I will set out a framework to help you get to grips with the core capabilities of AI and relate real business examples to each of these. I will provide approaches, methodologies and tools so that you can start your AI journey confidently and effectively.

Introducing the AI Framework

My AI Framework was developed through a need to be able to make sense of the plethora of information, misinformation and marketing-speak that is written and talked about in AI. I was not a computer coder or an AI developer, so I needed to put the world of AI into a language that business people like myself could understand. I was continually frustrated by the laziness in

using specific terminology in articles that were meant to help explain AI, and which only made the reader more confused than before. Terms like Artificial Intelligence, Intelligent Automation and Machine Learning were being used interchangeably despite being quite different things.

Through my work as a management consultant creating AI strategies for businesses, through reading many papers on the subject and speaking to other practitioners and experts, I managed to boil all the available information down into nine core capabilities for AI: Image Recognition, Speech Recognition, Natural Language Processing (NLP), Clustering, Natural Language Understanding (NLU), Optimisation, Prediction, Generation and Understanding. In theory, any AI application can be associated with one or more of these capabilities.

The first four of these all involve capturing information—getting structured data out of unstructured or big data. These Capture categories are the most mature today. There are many examples of each of these in everyday use: we encounter Speech Recognition when we use our smart speakers; we have Image Recognition automatically categorise our photographs; we have an NLP capability to read and categorise the emails we send, and we are categorised into like-minded Clusters when we buy something from an online retailer. AI efficiently captures all this unstructured and structured data that we give it and turns it into something useful (or intrusive, depending on your point of view, but that's a topic to be discussed in more detail later in the book).

The second group of NLU, Optimisation and Prediction are all trying to work out what is happening, often using the information that has just been captured. These also have applications in our daily lives. NLU works out what we want when we converse with a chatbot. The Optimisation capability (which includes problem-solving and planning as core elements) covers many uses, including working out the best route between your home and work. And then the Prediction capability tries to work out what will happen next—if we bought that book on early Japanese cinema, then we are likely to want to buy this other book on Akira Kurosawa. The Generation capability is a new one, added in since the last edition of this book. Previously, this was included as a minor adjunct to NLU, with only minimal usefulness. The advent of Transformers and Large Language Models in the last three years has changed this picture completely. We can now generate information in almost any format and field imaginable. We can ask about almost anything and generate an instant response; we can summarise huge corpora of information; we can generate novel images, music and video; we can generate and run computer code without being able to write code.

Once we get to Understanding, it's a different picture altogether. Understanding why something is happening really requires cognition; it requires many inputs, the ability to draw on many experiences and to conceptualise these into models that can be applied to different scenarios and uses. This is something that the human brain is extremely good at, but AI, to date, can't do with any level of understanding. Although Generative AI muddies the waters a bit, all of the previous examples of AI capabilities have specific uses (these are usually termed Narrow AI), but Understanding requires general AI, and this doesn't exist yet outside of our brains. Artificial General Intelligence, as it is known, is the holy grail of AI researchers, but it is still very theoretical at this stage. Even Generative AI, which has many use cases and can seem to suggest intelligence in its outputs, does not have any level of consciousness—at its core, it is an algorithm that effectively predicts the next words in a sentence. I will discuss the future of AI in the concluding chapter, but this book, as a practical guide to AI in business today, will inherently focus on those Narrow AI capabilities that can be implemented now.

You may already be starting to realise from some of the examples I have given that when AI is used in business, it is usually implemented as a combination of these individual capabilities strung together. Once the individual capabilities are understood, they can be combined to create meaningful solutions to business problems and challenges. For example, I could ring up a bank to ask for a loan: I could end up speaking to a machine rather than a human, in which case AI will first be turning my voice into individual words (Speech Recognition), working out what it is I want (NLU), deciding whether I should get the loan (Prediction) and then asking me whether I wanted to know more about car insurance because people like me tend to need loans to buy cars (Clustering). That's a fairly involved process that draws on key AI capabilities and doesn't have to involve a human being. The customer gets great service (the service is available day and night, the phone is answered straight away, and they get an immediate response to their query), the process is efficient and effective for the business (operating costs are low, the decision-making is consistent) and revenue is potentially increased (cross-selling additional products).

The AI Framework gives us a foundation to help understand what AI can do (and to cut through that marketing hype), but also to help us apply it to real business challenges. With this knowledge, we can answer questions such as: How will AI help me enhance customer service? How will it make my business processes more efficient? And how will it help me make better decisions? All of these are questions that I will explore in detail in the course of this book.

Defining AI

Interestingly, in most of the examples I have given so far, people often don't even realise they are dealing with AI. Some of the uses today, such as planning a route in our satnav or getting a phrase translated in our browser, are so ubiquitous that we forget that some really clever stuff is happening in the background. This has given rise to some tongue-in-cheek definitions of what AI is: some say it is anything that will happen in 20 years' time, and others that it is only AI when it looks like it does in the movies. But, for a book on AI, we do need a concise definition to work from.

The most useful definition of AI I have found is, unsurprisingly, from the Oxford English Dictionary, which states that AI is "the theory and development of computer systems able to perform tasks normally requiring human intelligence". This definition is a little bit circular since it includes the word 'intelligence', and that just raises the question of what intelligence is, but we won't be going into that philosophical debate here.

Another definition of AI that can be quite useful is from Andrew Ng, who was a co-founder and head of Google Brain, the former Chief Scientist at Baidu, and a bit of a rock star in the world of AI. He reckons that any cognitive process that takes a human under one second to process is a potential candidate for AI. Now, as the technologies get better and better, this number may increase over time, but for now, it gives us a useful benchmark for the capabilities of AI.

Another way to look at AI goes back to the very beginnings of the technology and a fundamental question: should these very clever technologies seek to replace the work humans are doing, or should they augment it? There is a famous story of two of the 'godfathers' of AI who were at the Massachusetts Institute of Technology (MIT): Marvin Minsky and Douglas Engelbart. Following a breakthrough in his research, Minsky declared, "We're going to make machines intelligent. We are going to make them conscious!" To which Engelbart reportedly replied: "You're going to do all that for the machines? What are you going to do for the people?" This debate is still raging today and is responsible for some of those 'robots will take over the world' headlines that I discussed at the top of this chapter. This book focuses as much as possible on the augmentation side of the debate, where the most value can be added, and it is this aspect that is explored further in a new chapter of this book, called 'Living with AI'.

A Brief Technology Overview

The technology behind AI is fiendishly clever. At its heart, there are algorithms: an algorithm is just a sequence of instructions or rules followed to complete a task, so it could simply be a recipe or a railway timetable. The algorithms that power AI are essentially very complicated statistical models—these models use probability to help find the best output from a range of inputs, sometimes with a specific goal attached (‘if a customer has watched these films, what other films would they also probably like to watch?’). This book is certainly not about explaining the underlying AI technology; in fact, it is deliberately void of technology jargon, but it is worth explaining some of the principles that underpin the technology.

Most AI today is a form of Machine Learning (ML). As the name suggests, the algorithm learns from the data it is given to build the model rather than the model being explicitly built by a person. This means that ML depends on having lots of data to learn from.

Machine Learning can be sub-categorised between ‘supervised’ and ‘unsupervised’ learning. Supervised learning is the most common approach and refers to situations where the AI system is trained using large amounts of labelled data. For example, if you wanted to have an AI that could identify pictures of dogs, then you would show it thousands of pictures, some of which had dogs and some of which didn’t. Crucially, all the pictures would have been labelled ‘a dog picture’ or ‘not a dog picture’. Using machine learning and all the training data, the system learns the inherent characteristics of what a dog looks like. It can then be tested on another set of similar data, which has also been tagged, but the tags haven’t been revealed to the system this time. If trained well enough, the system can identify the dogs in the pictures and correctly identify pictures with no dogs. It can then be let loose on real examples. And, if the people using your new ‘Is There a Dog in My Picture?’ app can feed back when it gets it right or not, then the system will continue to learn as it is being used.

With Unsupervised Learning, the system starts with a very large data set that will mean nothing to it. What the AI can do, though, is to spot clusters of similar points within the data. The AI is blissfully naive about what the data means; all it does is look for patterns and correlations in vast quantities of numbers. The great thing about this approach is that the user can also be naive—they don’t need to know what they are looking for or what the connections are—the AI does all that work. Once the clusters have been identified, predictions can be made for new inputs.

So, as an example, we may want to be able to work out the value of a house in a particular neighbourhood. The price of a house depends on many variables such as location, number of rooms, number of bathrooms, age, size of garden and so on, all of which make it difficult to predict its value. But, surely, there must be some complicated connection between all of these variables if only we could work it out. And that's exactly what the AI does. If it is fed enough base data, with each of those variables and the actual price, then it uses complex statistical analysis to find all the linear and non-linear connections—some variables may be very strong influencers on price, whilst others may be completely irrelevant. You can then input the same variables for a house where the price is unknown, and it can predict that value. The input data is structured, but the model created is a black box. This apparent lack of transparency is one of AI's Achilles's heels, but one that can be managed, which I'll discuss later in the book.

As well as the above two types of training, there are various other terms associated with AI, which I'll cover briefly here, although, for business executives, they only need to be understood at a superficial level. 'Neural Networks' is the term used to describe a certain category of ML algorithm where AI mimics the way that the brain processes information—many 'neurons' (approximately 86 billion in the case of the brain) are connected in various degrees of strength across a number of layers, and the strength of the connection can vary as the brain/machine learns. Complex neural networks, with many layers, give rise to Deep Learning, although the concepts are very much the same.

Generative AI exploits a number of different algorithmic approaches, but it is essentially the use of unsupervised learning trained across extremely large datasets. In the case of Large Language Models (such as ChatGPT), training data can be considered the whole internet (warts and all). There is then some refinement of the models using human feedback to score the responses. The model then predicts the next 'tokens' (parts of words) to display so that it makes grammatical and, as far as possible, factual sense.

There are obviously many more terms that are in common use in the AI world, including Decision Trees, Generative Adversarial Networks, reinforcement learning and Bayesian networks, but I will cover these only when absolutely necessary. As you will hopefully appreciate, this book focuses on the business application of AI rather than its underlying technologies.

About This Book

My working experience has been as a management consultant, helping organisations cope with the challenges of the time, from productivity improvement through change management and transformation to outsourcing and robotic process automation, and now AI. I first came across AI properly in my work in 2001. I was working as Chief Technology Officer in the Corporate Venturing unit of a global insurance provider—my role was to identify new technologies that we could invest in and bring into the firm as a foundation client (it was what we used to call the ‘incubator’ model). One of the technologies we invested in was based on the idea of ‘smart agents’ that could be used as an optimisation engine—each agent would have a specific goal and ‘negotiate’ with the other agents to determine the ideal collective outcome. So, for example, the system could determine the most effective way for trucks to pass through a port or the best way to generate the most revenue from the size and arrangement of newspaper advertisements. Although we didn’t call it AI at the time, this is effectively what it was—using computer algorithms to find optimal solutions to real problems.

Fast forward to 2024, and my work is focused almost exclusively on AI. I work with enterprises to help create their AI strategy—identifying opportunities for AI, finding the right solutions and creating the roadmap for implementation. I do this as someone who understands how the capabilities of AI might address business challenges and opportunities, as well as someone who can build solutions with computer code. I see myself as the ‘translator’ between technologists and the business. And with AI, the challenge of translating the technology is orders of magnitude greater than with traditional IT. This is why I wanted to write this book—to bring that understanding to where it can be used the best: on the front line of business.

So, this is not a book about the theoretical impact of AI and robots in 10 or 20 years’ time; neither is it a book about how to develop AI models. This is a book for practitioners of AI, people who want to use AI to make their businesses more competitive, innovative and future-proofed. That will happen only if the business leaders and executives understand what the capabilities of the technology are and how it can be applied in a practical way. That is the mission of this book: be as informed as possible about AI, but without getting dragged down by the technology, so that you can make the best decisions for your business. And it is also a heartfelt appeal: whatever you read or hear about AI, don’t believe the hype.