

Future of Business and Finance

Mayank Kejriwal

Artificial Intelligence for Industries of the Future

Beyond Facebook, Amazon, Microsoft
and Google



Springer

Future of Business and Finance

The Future of Business and Finance book series features professional works aimed at defining, describing and charting the future trends in these fields. The focus is mainly on strategic directions, technological advances, challenges and solutions which may affect the way we do business tomorrow, including the future of sustainability and governance practices. Mainly written by practitioners, consultants and academic thinkers, the books are intended to spark and inform further discussions and developments.

Mayank Kejriwal

Artificial Intelligence for Industries of the Future

Beyond Facebook, Amazon, Microsoft
and Google



Springer

Mayank Kejriwal
Information Sciences Institute, Ste 1001
University of Southern California
Marina Del Rey, CA, USA

ISSN 2662-2467 ISSN 2662-2475 (electronic)
Future of Business and Finance
ISBN 978-3-031-19038-4 ISBN 978-3-031-19039-1 (eBook)
<https://doi.org/10.1007/978-3-031-19039-1>

© The Editor(s) (if applicable) and The Author(s), under exclusive license to Springer Nature Switzerland AG 2023

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.

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

*To the future generation of my family: my
niece, Mishka, and my nephew, Atharv.*

Preface

Everyone has a view on artificial intelligence (AI) these days: scientists, policy makers, lawyers, engineers, journalists, consultants, and even artists. For a relatively young researcher in the field, this may seem like a welcome sign: after all, isn't public and government interest a sure-shot sign, as any, of increased funding and opportunities?

While this excitement can be heady, it has the unfortunate consequence of subjecting the field and its advancement to a degree of speculation that is not always supported by the evidence. I do not dispute that speculation, and its more creative cousin, imagination, can have intrinsic value; they prepare us for what might lie ahead, and they can be sources of inspiration for new ideas and modes of thought. At the same time, for both industry and government to make sound decisions and rationally allocate resources, it is important to be able to separate fact-based trends from pure speculation, and to also think about what the most reasonable outcomes might be, even if only the present state of the technology was taken into account.

The goal of this book is to set out the role of AI in “industries of the future.” It is not to speculate on “what” those industries might be, since history has taught us to be humble about our abilities to predict inventions in the long-term future. Who could reasonably have predicted, even in 2000, that a company like Apple would become a trillion-dollar (and counting) future smartphone maker or that commercial space tourism, pioneered by a start up, would soon be on the horizon? Not to mention the normalization of remote work, uniquely twenty-first century problems such as social media addiction, a decentralized cryptocurrency (Bitcoin) that takes the energy of a small country to “mine,” and an AI-composed artwork that has spurred international debate on whether an AI can be granted copyright on its work (Chap. 6)? In hindsight, it seems that we may have better luck (which is to say, very little) predicting business cycles than in predicting the path, let alone the impact, of advanced technology.

I would hope that my aim is less ambitious, but no less useful. Put simply, I am not going to spend much time speculating about *what* the inventions or industries of the future will be. Nor will I be discussing the projected impacts of technology that *currently* does not exist (and may not exist for another twenty years, or, for all we know, could be mainstream only a few years from now) might have. Instead, this book is predicated on the assumption that AI, as it exists today, already allows us to say something important about potential industries of the future. In other words, we

do not have to (necessarily) know what future AI will look like, any more than we need to try and define markets that do not exist yet, to still be able to *meaningfully* speak about what industry today should be thinking about as they set their sights toward an uncertain future.

What might some such features be? First, press and media hype aside, the sobering reality is that many companies are still a long way off from implementing AI today. In part, one issue is that measuring the return on investment (ROI) of AI is challenging, if not methodologically unprecedented. In Chap. 2, I go deeper into this issue after first introducing AI in Chap. 1, clarifying some of the terminological choices that I made when planning and writing this book. Fortunately, industry watchers and consultants have already thought about the thorny issue of AI implementation and ROI in some detail, guided undoubtedly by the operational, strategic, and pragmatic concerns and challenges that they have personally encountered in their conversations with clients and company leaders. My efforts in that chapter are to synthesize their lessons in a way that can be availed of (at least as a starting point) by any company looking to implement AI more vigorously throughout the fabric of their organization.

Even if we can't predict with great specificity what industries of the future might be, we can reasonably infer where they may emerge from. I argue that there are three important (and rather unsurprising) such sources: Big Tech, start-ups, and major corporations that are not Big Tech, but that are looking heavily to emerging technologies to improve their productivity and future growth prospects. In Chap. 3, I consider all three groups, and show that, in looking at the current advances being made in each, it is possible to say something (albeit with caution) about new AI-driven products and services that are already on the horizon, and could be impressive money-makers for the companies furthest along in developing and commercializing them. One such field of AI that I focus on as a mini-case study is Natural Language Processing (NLP) where truly remarkable progress has been achieved in both industry and academia over the last half decade. While much of the focus is on the United States, I also briefly discuss the important role of the Chinese Big Tech companies, many of which are also at the bleeding-edge of research in high-impact areas like Web search, NLP, and computer vision.

Irrespective of who the winners and losers will be in this race, whether nationally or internationally, there is no question anymore that both the government and human resources divisions within corporations will have to be involved. Chapter 4 considers how the workforce will be affected by “augmented” AI, which (in contrast with “automated” AI) assumes that AI is an advanced tool or implement that must work *alongside* humans to solve complex problems. Most likely, this is the variety of AI that will get implemented in practice in large organizations. This is not to say that every job or occupation will be equally affected. Similar to effects of technologies in the past (including the personal computer and word processing software), some jobs will get automated away, and additionally, some sectors will be more heavily affected than others. Similar to Chap. 3, I aim to use actual examples and consulting studies to show what these effects are likely to be. I present a small case study of a medical sub-field (radiology) that we may not think is at risk of upheaval by AI; in

fact, I draw on scholarly writing by actual radiologists to show that the changes are already happening, and radiologists need to adapt to the reality that AI is already quite good at (if not better than) some of the core tasks that radiologists are called upon to do today.

Chapter 5 dives deeper into AI ethics and policy. This is a rapidly evolving area, with each year bringing new developments to the fore. However, some recently passed laws and trends that are relevant to both AI and privacy (such as the European General Data Protection Regulation or the GDPR) are likely to stay, and we are already starting to see some of their effects. No company with a sufficiently large market size, or technological footprint, can afford to ignore these regulations and trends.

Chapter 6, wherein I conclude the book, is more future-facing than the others. I provide some guidance on where the future is headed. Once again, I try to avoid speculation by tying each issue that I discuss into ongoing events that provide an evidence-based foundation for why I included the issue (in that chapter) in the first place. One example (also mentioned earlier) is AI copyright, which, as futuristic as it might sound, is already being debated in patent and copyright agencies in several countries as part of a global project currently being led by a law graduate.

At this juncture, I would like to clarify my use of single quotes and double quotes. The former is used when I am drawing attention to a term, the usage of which may be less obvious or more questionable than it first seems, while the latter is always used when quoting directly from a source. Occasionally, I also used italicization and bold text to draw attention to a phrase, term, or its contextual usage.

To conclude, we are at an exciting juncture in our technological journey as a species. We may not know what is to come, or what specific form it will take, but current AI technology has advanced far enough as to enable us to make some reasonable claims about industries of the future. Although short, this book draws on a mix of scholarly sources, media reporting, and consulting pieces and commentary, to synthesize those claims in a single work. My hope is that it will prove to be useful to industry leaders and other interested stakeholders, both as an accessible review of contemporary perspectives on AI's forward-looking role in industry as well as a clarifying guide on the major issues that companies are likely to face as they commence on this exciting path.

Marina Del Rey, CA, USA
July 2022

Mayank Kejriwal

Acknowledgments

I have no hesitation in acknowledging that this book is built largely on the knowledge of others, but the most important contributor may be self-reporting by industry itself. In particular, a number of consulting firms have been investigating successful integration (or the possibility thereof) of AI in companies, including best practices and return on investment (ROI), and that knowledge has been disseminated regularly through whitepapers and reports. I have drawn upon that reporting extensively, with credit and citations accorded where due, for certain segments of this book, as their conclusions have been guided by actual (rather than academic or even hypothetical) problems encountered when working with clients. Certainly, without that dissemination, this book would not be possible. Credit also goes to the greater degree of openness (at least concerning AI research) with which many companies, both within Big Tech but also startups, are operating today. Open-source projects are now common in many industrial sectors, as is academic publishing and presentation of the best work of the researchers working in those companies. This openness allows the book to be non-speculative, and to focus on reality than on counterfactuals or hypotheticals.

My acknowledgments also extend to the growing body of commentators and journalists in the media who have increasingly started to report on AI-related issues, including political and ethical issues. Their coverage has informed many aspects of my own, and the book attempts to credit them by citing news reports where relevant. Some may argue that this makes the book less “scholarly” than a book that is more traditionally academic; however, I believe that, for a rapidly evolving topic that is focused on industry as a first-class citizen, reporting from reputed news sources as well as discussions in industry whitepapers do constitute a primary source of material.

Last, but not least, I credit policy research done (and disseminated) by governments and policy makers that have informed certain chapters in the book. Such research is usually conducted by advisory bodies and special councils comprising experts who often (voluntarily, and for free) lend their knowledge to advising governments and lawmakers on best practices for regulating and adopting AI. Their thought leadership informs this book, which is, ultimately, a mere synthesis and synopsis of their more extensive and original explorations.

Contents

1	Artificial Intelligence: An Introduction	1
1.1	Introduction	1
1.2	Artificial Intelligence (AI)	3
1.3	AI, Machine Learning, and Deep Learning	6
1.3.1	Types of Machine Learning	7
1.4	Industry 4.0 Versus Industries of the Future	10
1.5	Other (Non-AI) Drivers of Industries of the Future	10
1.5.1	Quantum Information Science (QIS)	11
1.5.2	5G and Advanced Communication	12
1.5.3	Advanced Manufacturing	13
1.5.4	Biotechnology	14
1.6	Where Will Industries of the Future Come From?	15
1.7	The Role of Research	17
1.8	Future Developments	21
	References	23
2	AI in Practice and Implementation: Issues and Costs	25
2.1	Introduction	25
2.2	Challenges in Implementing AI	26
2.2.1	Data Acquisition	27
2.2.2	Data Quality	29
2.2.3	Privacy and Compliance	30
2.2.4	AI Quality Metrics	31
2.3	Guidelines and Practices for Measuring Return on Investment (ROI) of AI Projects	34
2.3.1	Traditional Valuation Approaches and Their Pitfalls for Valuing AI Projects	35
2.3.2	Soft Versus Hard Returns and Investments	36
2.4	Digital Technology and the Productivity Puzzle	41
2.5	Conclusion	42
	References	42
3	AI in Industry Today	47
3.1	Introduction	47
3.2	AI in Big Tech	48
		xiii

3.2.1	Alphabet	49
3.2.2	Amazon	51
3.2.3	Meta	54
3.2.4	Other Big Tech: Microsoft and Apple	55
3.2.5	Other Large Tech Firms in the United States	55
3.2.6	The Chinese “Big Tech”	57
3.3	Large Firms Outside Big Tech	59
3.4	Startups and Small/Medium-Sized Enterprises (SBEs)	61
3.5	Case Study: Neural Language Models	63
3.5.1	Can Transformers Automate Software Engineers?	65
3.5.2	Applications Beyond NLP	66
3.5.3	Potential Ethical Concerns	66
3.5.4	Summary	67
3.6	Conclusion	68
	References	69
4	Augmented Artificial Intelligence	75
4.1	Introduction	75
4.2	Augmented AI Versus Complete Automation	76
4.3	Key Features and Example Applications	78
4.4	A Case Study in Augmented AI: Radiology	80
4.5	Changes in the Workforce	85
4.5.1	How Will Organizations Change?	86
4.5.2	Demand for Technological Skills	88
4.5.3	Cognitive Skills and the Future of Work: Is There a Mismatch?	89
4.5.4	New-Collar Versus White-Collar Jobs	90
4.5.5	Adaptation in the C-Suite	92
4.6	Automation and the Future of Work: Examples from Three Industrial Sectors	93
4.6.1	Banking and Insurance	94
4.6.2	Manufacturing	95
4.6.3	Retail	95
4.7	Conclusion	96
	References	97
5	AI Ethics and Policy	101
5.1	Introduction	101
5.2	AI Versus Digital Ethics	102
5.3	The Philosophy of Ethics: A Brief Review	104
5.4	AI Ethics in Policy	105
5.4.1	Case Study 1: The European Union General Data Protection Regulation (GDPR)	111
5.4.2	Case Study 2: The United States National Defense Authorization Act (NDAA)	114
5.5	AI Ethics in Research and Higher Education	116

5.6 Conclusion	117
References	118
6 What Is on the Horizon?	119
6.1 Introduction	119
6.2 Can AI Copyright Its Own Art?	120
6.3 Legal Issues Around Deepfakes	121
6.4 AI’s Explainability Crisis	123
6.5 More Vigorous Algorithmic Regulation	125
6.6 Increasing Convergence of Emerging Technologies	127
6.7 Concluding Notes	129
References	130
Glossary	137
References	141
Index	143

Acronyms

AGI	Artificial General Intelligence
AI	Artificial Intelligence
AI2	Allen Institute for Artificial Intelligence
AICT	Artificial Intelligence Capabilities and Transparency Act
AIM	Artificial Intelligence for the Military Act
API	Application Programming Interface
AWS	Amazon Web Services
BATX	Baidu, Alibaba, Tencent, and Xiaomi
BERT	Bidirectional Encoder Representations from Transformers
CASMI	Center for Advancing Safety of Machine Intelligence
CASP	Critical Assessment of Techniques for Protein Structure Prediction
CoE	Council of Europe
CSR	Common Sense Reasoning
DABUS	Device for the Autonomous Bootstrapping of Unified Sentience
DARPA	Defense Advanced Research Projects Agency
DL	Deep Learning
DoD	Department of Defense
DQD	Data Quality Dimension
EBITDA	Earnings Before Interest, Tax, Depreciation and Amortization
EHR	Electronic Health Record
EU	European Union
EY	Ernst & Young
FANG	Facebook, Amazon, Netflix, and Google
FAANG	Facebook, Amazon, Apple, Netflix, and Google
FDA	Food and Drug Administration
FTC	Federal Trade Commission
GAN	Generative Adversarial Network
GDP	Gross Domestic Product
GDPR	General Data Protection Regulation
GovAI	Centre for the Governance of AI
GPAI	Global Partnership on Artificial Intelligence
GPT	Generative Pre-Trained Transformer
GPU	Graphical Processing Units
HAI	Human-Centered AI

HEART	Hydrogen Electric and Automated Regional Transportation
HR	Human Resources
IDE	Integrated Development Environment
IE	Information Extraction
i.i.d	independent and identically distributed
IoT	Internet of Things
IP	Intellectual Property
IT	Information Technology
JEDI	Joint Enterprise Defense Infrastructure
KB	Knowledge Base
KG	Knowledge Graph
KPMG	Klynveld Peat Marwick Goerdele
LM	Language Model
LRM	Language Representation Model
LORELEI	Low Resource Languages for Emergent Incidents
LwLL	Learning with Less Labeling
M&A	Mergers & Acquisitions
MAMAA	Microsoft, Apple, Meta, Alphabet, and Amazon
MCS	Machine Common Sense
ML	Machine Learning
MOOC	Massive Open Online Course
MRI	Magnetic Resonance Imaging
NDAA	National Defense Authorization Act
NIST	National Institute of Standards and Technology
NLP	Natural Language Processing
NLU	Natural Language Understanding
NPV	Net Present Value
NPR	National Public Radio
NSCAI	National Security Commission on Artificial Intelligence
NSF	National Science Foundation
NYT	The New York Times
NYU	New York University
OECD	Organisation for Economic Co-operation and Development
PaLM	Pathways Language Model
PDSS	Personalized Decision Support Systems
PwC	PricewaterhouseCoopers
QIS	Quantum Information Science
R&D	Research & Development
RDBMS	Relational Database Management System
RL	Reinforcement Learning
ROI	Return on Investment
RPA	Robotic Process Automation
SAIL-ON	Science of Artificial Intelligence and Learning for Open-world Novelty
SBIR	Small Business Innovation Research

SEO	Search Engine Optimization
SME	Small and Medium-Sized Enterprise
STEM	Science, Technology, Engineering, and Mathematics
STTR	Small Business Technology Transfer
SVM	Support Vector Machine
TFQ	TensorFlow Quantum
TPU	Tensor Processing Units
UK	United Kingdom
UN	United Nations
UNESCO	United Nations Educational, Scientific and Cultural Organization
UNICRI	United Nations Interregional Crime and Justice Research Institute
US	United States
USA	United States of America
USPTO	United States Patent and Trademark Office
VLSI	Very-Large Scale Integration
WIPO	World Intellectual Property Organization
WSJ	The Wall Street Journal