Banking on (Artificial) Intelligence

Navigating the Realities of AI in Financial Services

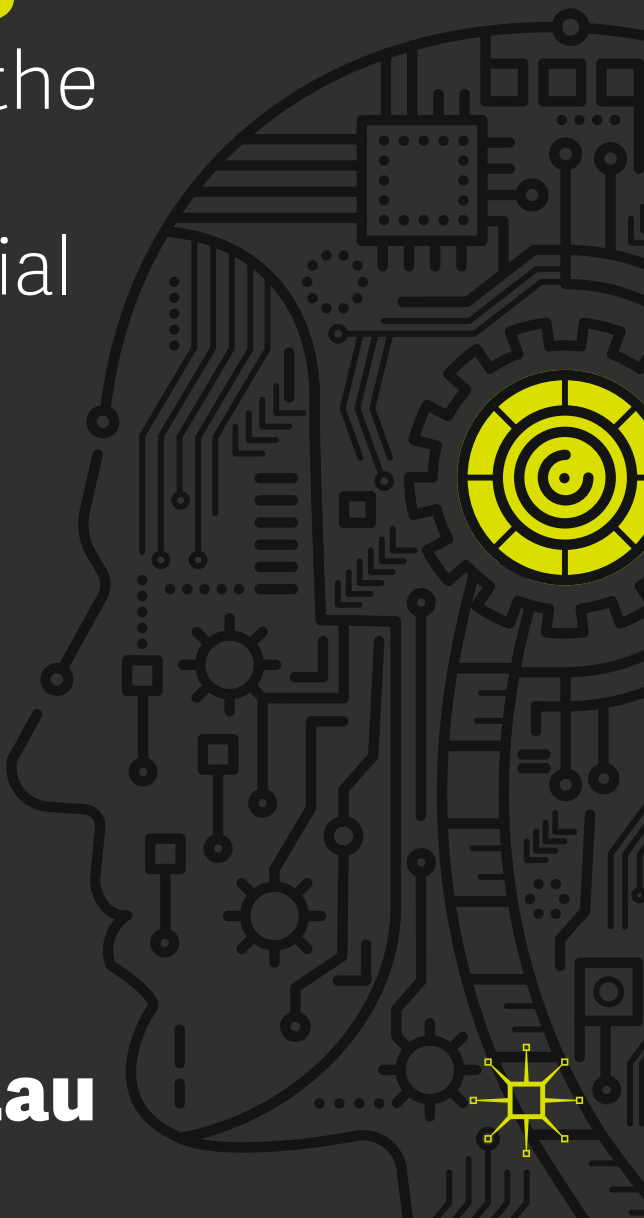

Theodora Lau

Banking on (Artificial) Intelligence

Theodora Lau

Banking on (Artificial) Intelligence

Navigating the Realities of AI in Financial Services

palgrave
macmillan

Theodora Lau
Unconventional Ventures
Fairfax, VA, USA

ISBN 978-3-031-81646-8 ISBN 978-3-031-81647-5 (eBook)
https://doi.org/10.1007/978-3-031-81647-5

Cover illustration: Lemberg Vector Studio

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

This book is dedicated to everyone who believe in a fairer and more equitable world for all.

To the Dreamers. The Doers. The Humans of AI.

Foreword

The current era of artificial intelligence expansion in financial services marks a pivotal moment in the trustworthiness of global economic systems. As we stand at this crossroads, the book you hold in your hands serves as an indispensable guide to navigating the complex landscape ahead.

The expansion of AI into finance is not merely a technological shift; it represents a fundamental reimagining of how we interact with money, credit, and economic opportunity. From algorithmic trading to personalized financial advice, to fraud detection and credit decisioning, AI is reshaping the financial sector at breakneck speed. This transformation promises immense benefits, but it also harbors significant risks that demand our collective attention, and importantly, our informed action.

As someone who studies the intersection of technology and society, I cannot overstate the importance of widespread AI literacy in this new era. Understanding the capabilities and limitations of AI is no longer a luxury reserved for technologists and data scientists—it is a necessity for everyone touched by the financial system. This book plays a crucial role in democratizing that knowledge, offering clear, accessible insights into the workings of AI in finance.

Theo has masterfully structured this work for us to have a comprehensive understanding of AI in financial services. Beginning with the historical context and current state of AI, the book progresses through practical applications, delving deep into the opportunities and challenges that lie ahead. What

sets this work apart is its unflinching examination of the ethical, societal, and environmental implications of AI in finance.

Particularly commendable is the emphasis on responsible AI development. In an age where algorithms can make decisions that profoundly impact our financial lives, issues of bias, transparency, and accountability are not merely academic concerns, but urgent practical matters. Theo rightly highlights, ensuring AI serves the interests of all—not just the privileged few—requires collective action, and that starts with informed citizens.

The exploration of AI's environmental impact is both timely and necessary. As we grapple with the climate crisis, we must confront the resource-intensive nature of AI development and deployment. This book offers valuable perspectives on balancing technological progress with environmental stewardship—a crucial consideration often overlooked in discussions about AI.

Perhaps most importantly, Theo is unafraid to delve into the potential for AI to exacerbate societal inequities and distrust. By thoughtfully examining economic inequality and the future of work, she provides thought provoking guidance for creating a more inclusive and equitable AI-driven financial landscape.

This book is essential reading for anyone impacted by the financial services industry—which, in our interconnected world, is virtually all of us. It equips readers with the knowledge and critical thinking skills necessary to engage thoughtfully in the AI enabled financial services sector. More importantly, it empowers us to take collective action to shape an AI-driven future that is equitable, sustainable, and truly serves the needs of all.

Chicago, IL, USA Reggie Townsend

Reggie Townsend is the Vice President of the SAS Data Ethics Practice and member of the National Artificial Intelligence Advisory Committee.

Preface

Banking in (artificial) intelligence is the continuation of the Beyond Good journey—one that sits at the intersection of technology and financial services. It is a book of hope, written during times of great uncertainty—a time when hope seems to be elusive. It is also a book about choices—a portrait of our society undergoing transformation led by technological advances and macroeconomic factors.

I am old enough to remember the days when I used to go to the AAA office to get paper maps before taking road trips. I vividly recall the frustration whenever I got lost, desperately trying to navigate my way back. That seemed like distant memories, especially since the proliferation of GPS technology and smartphones. In the past few years, I have become fascinated by the idea of using artificial intelligence as our sidekick. Why can't we treat navigating our personal finance journey the same way, I wonder? Imagine the power of having a co-pilot in our pocket, one that understands not only what we want, but how the changing needs and aspirations of our family and loved ones impact our own financial planning and trajectory. Think of it as a trusted personal advisor who can guide us along the life stages as things change.

The world of data is a world of possibilities. There is power within the bits and bytes that we can harness for good. Surely, we will face challenges. And we are in the midst of another AI hype cycle where many questions loom large. AI as a technology has the potential to greatly change how we learn, how we work, and how we live. Business models will be upended, and new ideas will be formed. Early movers and those with vast resources will stand

to reap the benefits. And the divide between the haves and the have-nots will become greater if we are not careful. There are 7164 living languages spoken around the world today. If AI is truly going to be our future, what will happen to the many languages and cultures that will be left behind simply because they are not connected to the larger digital ecosystem and not included in the AI models? Who is not included in this future is just as important as who is included. Progress is not true progress at the expense of the well-being of our society and planet at large.

I am reminded of the dualities of AI that we face. Multiple things can be true at the same time, that AI can greatly benefit humankind, but it can also cause great destruction to the environment and social cohesion. Banking on (artificial) intelligence is about exploring the nuance of the impact and finding balance. It is about facing the constraints head on—from data, to talent, and resources. But most importantly, it is about shining the light on those who are working hard to dismantle the roadblocks—the stories that give us hope. As it is the case with any innovation, we can't progress unless we try. We learn through failures. We fall. We get back up and try again. The road ahead may seem treacherous, but we have simply come too far and worked too hard to give up.

We are just at the very beginning of a transformation and so much is yet unknown. This book is not meant to be prescriptive on the roads we must travel. Rather, it's meant to be the beginning of more conversations across our industry and in our society. I hope this can serve as a bridge to bring us closer together towards one common purpose: To go Beyond Good—to embrace the realities of AI in a data intelligence world and build a fairer and equitable future for all. What happens next lies with our collective actions today.

Fairfax, USA Theodora Lau

Acknowledgements

Truth is, I never set out to be a writer. In fact, I still do not myself to be one—at least not one of traditional sense. I think of myself as more of a storyteller—someone who explores and narrates. If your eyes are the windows to your souls, then words are reflections of your mind.

I am forever grateful for the team at Palgrave Macmillian, especially Tula Weis, who gave me the opportunity to tell my story. And this could not have happened without Bronwyn Geyer, who has been instrumental in my author journey.

I also want to give a huge shoutout to Reggie Townsend and Brian Lee for all that you do for our communities and for being the beacons of light in the midst of uncertainty. Your foreword and afterword brought me to tears.

Thank you to everyone who have taken the time to share your thoughts with me and guide me on my thinking. It is a privilege and an honor to share your stories.

Thank you to my tribe from around the world, for supporting me through the years and for listening to me without judgement. Especially Penny Crosman, Tiffani Montez, and Katie DeGraff, who have opened doors for me in so many ways.

And most importantly, to my family: I simply could not have done this without you. To the kiddos, thank you for teaching me to dream and to question the status quo. I hope that more adults can see the world through your eyes and learn from you.

To all the changemakers and guardians of AI, thank you.

Contents

List of Figures

1

Introduction

The future will be powered by data intelligence.

It is certainly hard if not impossible to ignore the excitement around artificial intelligence (AI) nowadays.

By the time you are reading this first chapter, this probably feels ancient—which speaks to how quickly the space is evolving. As I type, Apple has just wrapped up the 2024 Apple Worldwide Developers Conference, where they waltzed back into the ecosystem with their response to the generative AI buzz with Apple Intelligence and much-needed makeover for Siri. Just a month prior, OpenAI had their Spring Update with their latest GPT-4o, which is already leaps and bounds beyond the first version of ChatGPT that the team unveiled in late 2022 and sent the world into frenzy. At the same time, Google also announced a radical change to search at their Google I/O 2024 conference that would not only change the advertising model that has propelled the company to where it is today but will also bring about a different future in terms of what we will see and read. More on that in later chapter.

Despite the hype, however, AI is not new. In fact, the concept has been around since mid 1950s. And AI has been employed by incumbent financial institutions and established fintechs for quite a while. In this first chapter, we will explore the history of AI and how it has been applied in our daily lives.

T. Lau, *Banking on (Artificial) Intelligence*, https://doi.org/10.1007/978-3-031-81647-5_1

What is Artificial Intelligence?

The term, artificial intelligence, was coined by John McCarthy, an American scientist, in 1955.[1] In the simplest terms, artificial intelligence refers to a set of technologies that enables machines to perform complex tasks that previously require humans to do, all the while processing large amounts of data in a fraction of time it would take a human.

The definition of AI was updated by the OECD to reflect recent developments in generative AI. It now reads as follows: "An AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment."[2]

The prominent use of AI is best illustrated by the role that the technology plays in our modern lives. From the moment we wake up when we check our emails or news feed, to when we step out of the house and navigate through traffic, AI is working silently in the background or as a co-pilot in lockstep. As the technology continues to evolve, we'd likely see new applications and techniques.

This is an important moment not only for the banking industry, but for our society at large. Not because we should be worried about AI overtaking humans. But rather, because of the past and present challenges, such as biased data and cost of operations, that could impact the trajectory of where the technology will go, who it benefits, and whether it will bring us together as a society or further divide us.

A Brief History of AI

First, let's travel back in time to 1942.

An American writer named Isaac Asimov published a short story titled "Runaround" in the Astounding Fiction magazine in 1942 and laid out his Three Laws of Robotics[3], which were intended for robots to follow to prevent harm to their human creators. The laws stipulate that a robot may not injure a human nor allow a human to come to harm; that a robot must obey the orders given by humans and protect its own existence, so long as such action does not conflict with the prior law where the well-being of humans comes first.

Soon after, in 1950, British mathematician and computer scientist Alan Turing proposed an imitation game to determine whether machines can think. The game, now known as the Turing Test, involves having a remote human interrogator ask a computer and a human a set of questions, and the interrogator needs to be able to tell them apart based on responses they provide.

Turing went on to predict that in 50 years' time, computers would be able to play this imitation game so well that an average interrogator would not have more than a 70 percent chance of being able to tell machine apart from humans after five minutes of questioning.[4] Though I think it is safe to say that year 2000 came and went, and we have yet to see a machine that can *think* in the way Turing predicted.

The Golden Age of AI

Enthusiasm and hope followed the Turing Test prediction—during a period otherwise known as the golden age of AI. It started in 1956 when John McCarthy first coined the term artificial intelligence during a summer conference in Dartmouth University. Some of the advancements made during this period include:

- Shakey the robot—the first mobile robot to move around autonomously by sensing its surroundings. Developed by SRI (what was then Stanford Research Institute) from 1966 to 1972, Shakey's AI ecosystem was composed of multiple key components, including a TV camera and optical range finder (computer vision), an antenna radio link (communication system), bump detectors (navigational system), and a push bar to move objects.[5] Its layered software architecture was the first time it was used in robotics and the combination with the rest of the technology ecosystem made Shakey the model for future generations of AI-enabled robotic systems.
- SHRDLU—a program for understanding natural language. It was written by Terry Winograd at the MIT Artificial Intelligence Laboratory between 1968 and 1970 and could carry out a simple dialogue with a user to move colored blocks around on the table.[6] Its initial success helped inspire a series of efforts in commercializing AI.

In fact, the field was filled with so much excitement that in 1970, Marvin Minsky, one of the two founders of the MIT Computer Science and Artificial Intelligence Laboratory, predicted the imminent arrival of a machine with general intelligence of an average human in three to eight years.

The First AI Winter

After a series of disappointments resulting from failure in deploying AI in real-world scenarios, however, reality set in. Unlike modern-day AI, access to data was limited, cost of compute was high, and not to mention, there was simply not enough processing power to create or solve anything useful. We will examine some of these factors further in Chapter 4 of the book as cost remains one of the biggest constraints for the adoption and scaling of AI and will impact who has access to the technology and whom it will benefit.

As the saying goes, just because you can, doesn't mean that you should. For technology to be readily adopted, it needs to demonstrate true value and solve a real problem.

A historical moment came about in 1973 when Professor Sir James Lighthill of Cambridge University indicated that "the general-purpose robot is a mirage".[7]

The expectation that we could achieve human-level intelligence in such a short period of time was simply unrealistic. In fact, American philosopher, John Searle, presented the Chinese room argument in his paper, "Minds, Brains, and Programs," in 1980, to demonstrate that artificial intelligence is indeed artificial.[8] Just because a computer program can simulate intelligence by manipulating strings of Chinese characters does not equate to true understanding of the characters themselves as they do not bear any meanings for the machine.

Interests in AI and funding slowed down. The first AI Winter has arrived.

The AI Boom

But the slowdown did not last long. Beginning in 1980s, companies began to adopt expert systems—programs that mimicked the decision-making process of human experts, by leveraging knowledge from experts to answer questions or solve programs, using a series of if–then rules. With potential for commercial value, interests in AI returned, with focus on knowledge engineering.

XCON, also known as R1, was a rule-based system written by John McDermott of Carnegie Mellon University for Digital Equipment Corporation (DEC) to assist in the ordering of DEC's VAX-11/780 systems by ensuring that the proper components were configured in the order. The remarks that John made in his 1980 paper seemed to foretell the future commercial uses of AI: "The configurations that it produces are consistently adequate, and the information that it makes available to the technicians who physically assemble systems is far more detailed than that produced by the humans." (McDermott 1980).[9]

Another notable development came in 1984 when Douglas Lenat founded a project known as CYC, which he dubbed ontological engineering.[10] He argued that the only way for a computer program to have common sense is to import all the facts and knowledge into the program, so that it can process and read, and eventually, understand the meaning of human concepts. "Think of it as the tens of millions of rules of thumb about how the world works that are almost never explicitly communicated. Beyond these implicit rules, though, commonsense systems need to make proper deductions from them and from other, explicit statements" (Lenat 2020).[11]

By mid 1980s, another approach to AI also began to take shape. Started by Australian American scientist Rodney Brooks at the MIT AI Laboratory, nouvelle AI relies on its sensors and the outside world to learn, as opposed to the classical AI, which needs predefined rules and logic.[12]

The Second AI Winter

As the hype increased during the AI boom, the promises of what the technology could achieve did not materialize. The Fifth Generation Computer Project (FGCP), a $400 million 10-year endeavor started in 1982 by the Japanese government to fund expert systems and develop computers with reasoning capabilities, ended in 1992.[13] Unfortunately, the effort did not produce the breakthroughs that researchers had hoped, and it proved to be too ambitious and cost prohibitive. Expert systems required an immense of data to operate and storage was still expensive.

Commercial interest waned and funding decreased once again, marking a brief period of a second AI winter.