

College of Humanities
Master of Financial
Technology-
Fin Tech

Artificial Intelligence for Financial Technology

FNT 626

Lecture 01

Prepared By:

Dr. Abeir Moussa



Artificial Intelligence for Financial Technology

FNT 626- AI for Fin Tech

Lecture 01

LECTURE 01

CHAPTER ONE

FOUNDATIONS of ARTIFICIAL INTELLIGENCE and MACHINE LEARNING

LECTURE 01

Learning Objectives

1. Understand the Fundamentals of Artificial Intelligence (AI),
2. Understand the fundamental concepts of machine learning,
3. Explore the principles of neural networks and deep learning,
4. Gain insights into different machine learning algorithms and models,
5. Understand data preprocessing techniques, feature selection,
6. Explore AI Applications in Various Industries,
7. Address the ethical implications and regulatory challenges associated with the adoption of AI and machine learning,
8. Explore real-world case studies and practical applications of AI and machine learning.

LECTURE 01

Lecture Outline

1. INTRODUCTION
2. A BRIEF HISTORY OF AI
 - 2.1 The Early Years
 - 2.2 The AI Industry Flourishes
 - 2.3 The Current State of AI
3. THE AI LANDSCAPE
 - 3.1 The Scope of AI
 - 3.2 AI: Algorithmic Classes
4. APPLICATIONS IN THE FINANCIAL INDUSTRY
 - 4.1 Banking
 - 4.2 Investing
 - 4.3 Insurance
 - 4.4 AI's Challenges
5. CONCLUSION

Ch1: Foundations of artificial intelligence and machine learning**1. INTRODUCTION**

The notion of intelligence is at the heart of human nature, as reflected in the very name of our species – homo sapiens ('man the wise'). For thousands of years, we have sought to comprehend how 'a mere handful of matter can perceive, understand, predict and manipulate a world far larger and more complicated than itself'. The notion of intelligence cannot be fully captured in formulaic terms, which has led scientists to advance a plethora of divergent definitions. In Mainstream Science on Intelligence, a number of prominent scientists settled on the following characterization of intelligence: 'A very general mental capability that, among other things, involves the ability to reason, plan, solve problems, think abstractly, comprehend complex ideas, learn quickly and learn from experience.' To arrive at these capabilities, intelligence is said to integrate cognitive functions such as perception, attention, memory, language and planning. Ever since computers were invented, humans have also wondered whether machines might be able to exhibit intelligent behavior. Artificial intelligence

Ch1: Foundations of artificial intelligence and machine learning

1. INTRODUCTION

Foundations of artificial intelligence and machine learning (AI) is the interdisciplinary field that studies how machines can be made to think or act in an intelligent manner.

The broad foundations of AI include computer science, mathematics, neuroscience, psychology, linguistics and philosophy. In view of our abstract understanding of intelligence, it is perhaps unsurprising that the standard for technology that qualifies as AI is fuzzy and time-bound. In general, the AI label is attached to machines (that is, computers comprising hardware and software) that are able to solve complex problems, which may involve acting or thinking in a human-like manner.

In turn, machine learning (ML) is a subfield of AI that comprises algorithms that improve their performance with experience, as opposed to being programmed to follow predetermined rules. Through the lectures, several applications of AI in the financial industry will be highlighted, focusing on the banking, investing and insurance sectors.

2. A BRIEF HISTORY OF AI

2.1 The Early Years

2.2 The AI Industry Flourishes

2.3 The Current State of AI



LECTURE 01

Ch1: Foundations of artificial intelligence and machine learning

Ch1: Foundations of artificial intelligence and machine learning

2. A BRIEF HISTORY OF AI

2.1 The Early Years

Any meaningful introduction to AI ought to commence with a brief exploration of the field's historical milestones, which have shaped our expectations of machines' capabilities and their limitations.

Theory of computation demonstrated that any form of computation could be described digitally. Indeed, it was around this time that Pitts and McCulloch started to explore what researchers would later refer to as neural networks.

In 1950, Turing published a landmark paper in which he asked a short yet challenging question: 'Can machines think?' Turing recognized that the concept of 'thought' was elusive, opting to sidestep the ambiguity by advancing a question that he regarded to be closely related: 'Are there imaginable computers that would do well in the imitation game?' The imitation game, or Turing Test, involves an interrogator trying to find out which of two participants is a machine by asking questions and evaluating each participant's written responses.

Ch1: Foundations of artificial intelligence and machine learning**2. A BRIEF HISTORY OF AI cont.****2.1 The Early Years cont.**

The test's usefulness as a measure of AI's success was famously mocked by The Economist in an article entitled 'Artificial Stupidity', in which the author noted that human behavior is often remarkably unintelligent.¹⁴ As a result, AI researchers have not devoted much time to building machines with the sole purpose of passing the Turing Test.

Russell and Norvig drew an analogy with aeronautical engineers, who do not define the goal of their field as making 'machines that fly so exactly like pigeons that they can fool pigeons'.¹⁵ Nevertheless, Turing's paper remains relevant to this day, as it introduced fundamental concepts like machine learning and reinforcement learning, which are explored further below.

Ch1: Foundations of artificial intelligence and machine learning

2. A BRIEF HISTORY OF AI cont.

2.1 The Early Years cont.

In the summer of 1956, a small team of scientists led by John McCarthy organized a two-month conference at Dartmouth College to study the topic of 'artificial intelligence'.

This marked the first usage of the term, though Russell and Norvig note that perhaps the term 'computational rationality' would have been more precise and less threatening. The early years of AI were marked by great expectations and enthusiasm. Given the primitive computers and programming tools that were available at the time, and the fact that computers were perceived to be capable of little more than rudimentary arithmetic operations, people were astonished when-ever a computer did anything remotely clever. In his paper, Turing had drawn up a list of tasks that the intellectual establishment assumed machines would never be able to perform.

Ch1: Foundations of artificial intelligence and machine learning

LECTURE 01

2. A BRIEF HISTORY OF AI cont.

2.1 The Early Years cont.

Presented with this challenge, many AI researchers naturally responded by creating AI programs that surpassed one after the other. McCarthy would later refer to this period as the ‘Look, Ma, no hands!’ era. In 1958, McCarthy (then at MIT) defined the high-level language Lisp, which would become the main language for AI programs for the next 30 years. From the earliest days, AI researchers did not shy away from making bold predictions of the field's coming successes, which served as reference points for both industry insiders and outsiders. A period of slow development – commonly referred to as the first AI winter – ensued from 1974 to 1980.

Ch1: Foundations of artificial intelligence and machine learning

LECTURE 01

2. A BRIEF HISTORY OF AI cont.

2.2 The AI Industry Flourishes

The second AI wave took place in the 1980s, thanks to the adoption of 'expert systems' by major corporations around the world. Indeed, expert systems proved to be too costly to maintain and difficult to upgrade, as they were unable to learn from experience.

Kasparov stated that he felt 'a new kind of intelligence' across the board, one that could sense the virtue of sacrificing a decisive short-term advantage in favor of an ultimate victory. On the other hand, many critics were quick to dismiss the machine's achievement, arguing that the use of 'brute force', where the machine processes all possible moves to compute its best strategy on a probabilistic basis, was not representative of real intelligence.

Ch1: Foundations of artificial intelligence and machine learning

LECTURE 01

2. A BRIEF HISTORY OF AI cont.

2.2 The AI Industry Flourishes

Tesler argues that people perceive AI to be 'whatever machines haven't done yet' because humanity is defined by reference to our 'allegedly unique intelligence'.²¹ While the performance of advanced calculations in the 1950s may have marveled the bystander, most would now regard such a feat as banal and unexciting, rendering it unmeritorious of the AI label.

Ch1: Foundations of artificial intelligence and machine learning

2. A BRIEF HISTORY OF AI cont.

2.3 The Current State of AI

Since the turn of the century, the widespread adoption of the internet has led to an exponential increase in the generation of publicly accessible data. The terms 'big data' and 'data science' are relevant here, as they refer to the analysis (or 'mining') of large datasets, often by using AI to identify patterns and derive useful outputs from the data.

This growing supply of data has catalyzed the rate of progress in AI, specifically in the machine learning subfield, where the accuracy of a model is linked to the amount of high-quality data that is fed to it.

Ch1: Foundations of artificial intelligence and machine learning

LECTURE 01

2. A BRIEF HISTORY OF AI cont.

2.3 The Current State of AI

At the sometime, AI practitioners are no longer constrained by high costs or the need to run processes locally, as big tech companies like Amazon (Amazon Web Services), Microsoft (Azure) and Google (Google Cloud Platform) all offer competitively priced, cloud-based AI frameworks.

In addition, the raw processing power of computers has increased exponentially since 1993, roughly as anticipated by Moore's Law.

3. THE AI LANDSCAPE

3.1 The Scope of AI

3.2 AI: Algorithmic Classes



LECTURE 01

Ch1: Foundations of artificial intelligence and machine learning

Ch1: Foundations of artificial intelligence and machine learning

3. THE AI LANDSCAPE

3.1 The Scope of AI

At inception, the field of AI sought to replicate human thinking in machines. As such, researchers often categorize a machine's intellect based on how it compares with that of humans.

Even Watson, IBM's most advanced question-answering computing system, which benefits from a knowledge base that branches into multiple domains, is only able to resolve a fraction of the problems that our species has collectively conquered. **Artificial general intelligence (AGI)** refers to the intelligence of a machine that is able to successfully perform any intellectual task that a human can do, including learning, planning, decision-making under uncertainty, joking and perhaps even deceiving or manipulating others.

Ch1: Foundations of artificial intelligence and machine learning

3. THE AI LANDSCAPE cont.

3.1 The Scope of AI cont.

In a recent report, **the University of Oxford's** Future of Humanity Institute surveyed a panel of AI researchers on timelines for AGI's arrival and concluded that 'there is a 50% chance of AI outperforming humans in all tasks in 45 years'. However, the answers given by individual AI practitioners tend to range from 'in a few years' to 'never'. **Artificial superintelligence (ASI)** refers to a machine that greatly exceeds the cognitive performance of humans in virtually all domains of interest. It is sometimes argued that a machine that has developed AGI capabilities would be able to achieve ASI status in a short timeframe, perhaps even on the someday, as the iterations of recursive self-improvement accelerate.

Ch1: Foundations of artificial intelligence and machine learning

3. THE AI LANDSCAPE cont.

3.1 The Scope of AI cont.

Additionally, 'the first ultra intelligent machine is the last invention that man need ever make, provided that the machine is docile enough to tell us how to keep it under control'. Yudofsky proposes that research be undertaken to create **friendly AI (FAI)** by introducing checks and balances via mechanism design and programming utility functions that remain friendly as the AI system evolves.

Ch1: Foundations of artificial intelligence and machine learning

LECTURE 01

3. THE AI LANDSCAPE cont.

3.1 The Scope of AI cont.

In the absence of safeguards, Omohundro posits that an innocuous chess program with the ability to self-improve could pose an existential risk to humans. In turn, Maini and Sabri observe that 'it's impossible to predict what will happen, just as mice who intentionally designed a human might have trouble predicting what the human would do to their world.

A distinction is drawn between machines that merely simulate a mind (weak AI) and machines that actually have a mind (strong AI). To pass the Turing Test, weak AI would suffice, as it is only the machine's outputs that matter, as opposed to the process through which these are reached.

The extent to which this distinction can and should be investigated has been brought into question, with suggestions that this issues as relevant as the question of whether submarines can 'swim' or planes can 'fly'.

Ch1: Foundations of artificial intelligence and machine learning

3. THE AI LANDSCAPE cont.

3.2 AI: Algorithmic Classes

The field of AI can also be subdivided based on the types of algorithms that are used to create intelligent machines.

In the context of software, AI is a broad term that is associated with programs that succeed at performing a (complex) task. Due to the AI effect, many people no longer consider these achievements to be representative of 'real' intelligence. Nowadays, the term AI is becoming synonymous with **machine learning (ML)**. To achieve this, developers must typically select an adequate model for the task, transform the data into features that can be processed by the model (through feature engineering), and monitor the model's performance over time. The end goal is to produce an algorithm that has a low rate of error (perhaps lower than that of a human) and is able to improve over time. ML is particularly useful in situations where it would be sub-optimal to program these algorithms, for example if there is a need to eliminate human bias or fatigue-based error.

Ch1: Foundations of artificial intelligence and machine learning

3. THE AI LANDSCAPE cont.

3.2 AI: Algorithmic Classes cont.

For instance, Singh explains that the seemingly accurate algorithm that his students had trained to distinguish huskies from wolves was in fact just learning to identify snow in the images – ‘if snow: return wolf; else: return husky’. Research is being conducted on explainable AI (XAI), which aims to address the opacity arising from black box ML algorithms by describing in an intelligible manner the variables to which the algorithm attributes weight as it evolves.

By way of reference, estimates suggest that there are only 1080 atoms in the universe.

Ch1: Foundations of artificial intelligence and machine learning

LECTURE 01

3. THE AI LANDSCAPE cont.

3.2 AI: Algorithmic Classes cont.

There are three main methods that can be used to train ML algorithms. In supervised learning, the computer is fed a series of inputs and told what the corresponding output is (e.g., 'high' or 'low').

The objective of the MDP is to find the best policy (i.e., succession of actions for each state) to maximize its reward. If the agent gets closer to the optimal state, it receives feedback on its performance by way of the reward.

DL also draws from the learning methods that are used more broadly in ML. **DeepMind's AlphaGo** is an example of a software agent that uses a DNN to decide which action to undertake.

4. APPLICATIONS IN THE FINANCIAL INDUSTRY

4.1 Banking

4.2 Investing

4.3 Insurance

4.4 AI's Challenges



LECTURE 01

Ch1: Foundations of artificial intelligence and machine learning

4. APPLICATIONS IN THE FINANCIAL INDUSTRY

AI, and ML in particular, nowadays plays an integral role in many industries, ranging from healthcare to advertising and manufacturing. The financial industry is certainly no exception, with myriad AI applications being deployed across the banking, investing and insurance sectors. In this section, we high-light applications in each of these sectors and discuss their associated benefits.

4.1 Banking

As mentioned that AI algorithms can be used to produce credit scores. The goal in using these algorithms is therefore to arrive at a faster, less manual, and more accurate system for risk-based pricing. AI is also assisting banks' compliance departments, freeing them from manual and repetitive tasks.

Ch1: Foundations of artificial intelligence and machine learning

4. APPLICATIONS IN THE FINANCIAL INDUSTRY cont.

4.1 Banking cont.

AI-powered software can apply data mining techniques to extract relevant phrases from a sea of documents. An example of this is **JP Morgan's contract intelligence (COIN)** program, which in 2017 took over the job of reviewing approximately 12,000 new commercial loan contracts per year. The company claimed that the program replaced 360,000 hours of staff time between lawyers and loan officers. Since many of these contracts are standard in form, COIN was able to flag issues in certain documents in a matter of seconds. In addition, it appeared to have a lower error rate than humans.

Ch1: Foundations of artificial intelligence and machine learning

4. APPLICATIONS IN THE FINANCIAL INDUSTRY cont.

4.1 Banking cont.

Banks are also making use of **AI-powered chatbots** to assist with customers' everyday queries, provide account data on demand, and help customers to meet their saving goals. These chatbots are able to recognize specific 'trigger words' in customer queries and to draw an adequate response from a message pool. Through the use of chatbots, banks are able to offer 24/7 support, cut waiting times, and enable support staff to focus on higher-value tasks.

Ch1: Foundations of artificial intelligence and machine learning

LECTURE 01

4.APPLICATIONS IN THE FINANCIAL INDUSTRY cont.**4.2 Investing**

Machine learning is used extensively in the investment industry, particularly in the realm of quantitative and systematic investing. A novel application of ML in this space is **'nowcasting'**, which entails the real-time analysis of big data, including alternative (i.e., non-traditional) datasets. López de Prado argues that this practice can yield predictions that are more accurate than historical data (forecasts), on which modern financial theory is built.

In addition, he contends that with increasing competition, excess returns are becoming 'microscopic', meaning that traders will struggle to outperform in the long run if their analysis is limited to the capabilities of an Excel sheet.

Ch1: Foundations of artificial intelligence and machine learning

LECTURE 01

4. APPLICATIONS IN THE FINANCIAL INDUSTRY cont.

4.2 Investing cont.

For instance, social media platforms generate alternative data that can be analyzed with deep learning techniques in an attempt to infer retail sentiment for traded assets.

AI is also being used to assess investment opportunities in private markets. A somewhat exotic example of this is 'VITAL', a program which Deep Knowledge (a life sciences venture capital firm in Hong Kong) appointed as a voting member of its board of directors in 2014. The goal was for VITAL to analyze potential investments in biotechnology companies, which have a very high rate of failure.

Ch1: Foundations of artificial intelligence and machine learning

LECTURE 01

4.APPLICATIONS IN THE FINANCIAL INDUSTRY cont.

4.2 Investing cont.

In the wealth management sector, ‘robo-advisers’ are AI-driven algorithms that dynamically manage a portfolio of assets based on predefined parameters. The terminology here is somewhat misleading, as no robots (i.e., machines with the ability to exert physical actions) are involved in this process. These algorithms can select assets for inclusion in a personal portfolio on the basis of risk-return preferences and liquidity needs, as well as other investment factors (e.g., ESG considerations). The algorithms are able to place trades automatically, rebalance the portfolio, and generate regular performance reports.

Ch1: Foundations of artificial intelligence and machine learning

LECTURE 01

4. APPLICATIONS IN THE FINANCIAL INDUSTRY cont.

4.3 Insurance

AI-based applications are being used in the insurance industry to derive risk-based scores based on applicants' attributes and to price insurance policies. Companies like Lemonade seek to automate the entire process of obtaining a policy through a **chatbot** experience. The whole process can be completed in a matter of minutes, without intermediaries, and avoiding the need for paperwork.

Ch1: Foundations of artificial intelligence and machine learning

LECTURE 01

4. APPLICATIONS IN THE FINANCIAL INDUSTRY cont.

4.3 Insurance

AI applications are also being used to assist fraud-prevention processes. Deloitte estimates that fraud-related costs amount to 10% of insurance companies' overall claims expenditure. AI can assist insurers during the claim-handling process by querying the events asserted in an insurance claim. For instance, if a driver involved in a car accident states that adverse weather conditions impacted the car's braking distance, the AI can pick this up and check the weather reports to confirm whether this was indeed the case.

Ch1: Foundations of artificial intelligence and machine learning

LECTURE 01

4.APPLICATIONS IN THE FINANCIAL INDUSTRY cont.

4.3 Insurance cont.

AI can also help to boost insurance companies' operational efficiency in distributing payouts. By way of example, **in 2017 a Japanese life insurer (Fukoku Mutual Life Insurance Co)** introduced an AI application based **on IBM Watson** to accelerate the processing of medical claims. This application is able to calculate the applicable payout based on case-specific factors, such as the relevant policy, period of hospitalization, administered procedures and medical history.

To achieve this, the application amalgamates and analyses information from a variety of sources (including hospital records and internal files). The calculated payout is then sent to a member of staff for approval, before it is released to the policyholder.

Ch1: Foundations of artificial intelligence and machine learning**4.APPLICATIONS IN THE FINANCIAL INDUSTRY cont.****4.4 AI's Challenges**

AI is already playing a valuable role in the financial industry. Nonetheless, despite the advantages and opportunities associated with AI, there are several challenges that still need to be addressed. In addition to technical issues, there are legal, philosophical, governance, and ethical considerations that stem from the use of this technology.

Ch1: Foundations of artificial intelligence and machine learning

LECTURE 01

4.APPLICATIONS IN THE FINANCIAL INDUSTRY cont.

4.4 AI's Challenges

1. Data Privacy and Security:

The financial industry handles sensitive customer data, making data privacy and security a paramount concern. Ensuring that AI systems comply with stringent data protection regulations **such as GDPR and safeguard** against cyber threats is crucial.

2. Regulatory Compliance:

Adhering to complex and evolving regulatory requirements while integrating AI into financial processes. Compliance with regulations such as **KYC (Know Your Customer)**, **AML (Anti-Money Laundering)**, and financial reporting standards presents a significant challenge.

Ch1: Foundations of artificial intelligence and machine learning

4. APPLICATIONS IN THE FINANCIAL INDUSTRY cont.

4.4 AI's Challenges

3. Interpretability and Explainability:

The need for transparent and explainable AI models in financial decision-making. Regulators and stakeholders often require clear explanations for AI-driven decisions, especially in risk assessment, investment strategies, and credit underwriting.

4. Bias and Fairness:

Mitigating bias in AI algorithms, particularly in credit scoring, loan approvals, and risk assessment. Ensuring fairness and non-discrimination in AI-driven financial decisions is essential to maintain trust and regulatory compliance.

Ch1: Foundations of artificial intelligence and machine learning

4. APPLICATIONS IN THE FINANCIAL INDUSTRY cont.

4.4 AI's Challenges

5. Risk Management:

Developing AI models that accurately assess and manage financial risks, including market volatility, credit risk, and fraud detection. Ensuring the reliability and robustness of AI systems in risk management is critical.

6. Ethical Use of AI:

Addressing ethical considerations in the use of AI for financial services, including the responsible handling of customer data, ethical investment practices, and the ethical implications of algorithmic decision-making.

Ch1: Foundations of artificial intelligence and machine learning

4.APPLICATIONS IN THE FINANCIAL INDUSTRY cont.

4.4 AI's Challenges

7. Explainable AI in Trading:

Ensuring transparency and interpretability in algorithmic trading systems. The need to understand and explain the decision-making processes of AI-driven trading algorithms is crucial for compliance and risk management.

8. Customer Trust and Communication:

Building and maintaining customer trust in AI-driven financial services. Communicating the benefits and limitations of AI to customers and ensuring that AI enhances, rather than detracts from, the customer experience.

Ch1: Foundations of artificial intelligence and machine learning

4.APPLICATIONS IN THE FINANCIAL INDUSTRY cont.

4.4 AI's Challenges

9. Operational Resilience:

Ensuring the resilience and reliability of AI systems in the face of operational disruptions, cyber threats, and system failures. Maintaining operational continuity and security is essential in AI adoption.

10. Talent and Expertise:

Acquiring and retaining talent with expertise in both finance and AI. The demand for professionals who understand the intersection of finance, data science, and AI presents a talent acquisition challenge for financial institutions.

Ch1: Foundations of artificial intelligence and machine learning

LECTURE 01

5. CONCLUSION

Artificial intelligence (AI) is an interdisciplinary field of study that focuses on building machines that are able to think and, in particular, act in an intelligent manner.

In turn, deep learning (DL) is a subfield of ML that uses deep neural networks (DNN) to identify useful patterns in the input data. On the technical side, this chapter draws a distinction between artificial narrow intelligence (ANI), artificial general intelligence (AGI) and artificial superintelligence (ASI).

Also, the distinction between weak AI and strong AI based on consciousness have been explored, though concerns have been raised about the relevance and falsifiability of these concepts. AI is already playing an important role in the financial industry, including in the banking, investing and insurance sectors.

Artificial Intelligence for Financial Technology

References:

Required Reference:

Remolina, N. and Gurrea-Martinez, A. eds., 2023. Artificial intelligence in finance: Challenges, opportunities and regulatory developments. Edward Elgar.

Recommended Readings

Chishti, S., 2020. The AI book: the artificial intelligence handbook for investors, entrepreneurs and fintech visionaries. John Wiley & Sons.



The End of Lecture

Thank You