

Financial Risk Forecasting

Chapter 12


Machine Learning and Artificial Intelligence

Jon Danielsson ©2024
London School of Economics

To accompany
Financial Risk Forecasting
FinancialRiskForecasting.com

Published by Wiley 2011

Version 9.0, August 2024



Authority use of AI

Central banks, various regulators, ...

- The authorities are much slower in engaging with AI than the private sector
- Most appear to ignore AI entirely
- The risk is that they will be faced with *fait accompli*
- AI is entering by stealth into the authorities
- The authorities will have no choice but to keep up with AI if they wish to be relevant
- We return to this theme later

What is artificial intelligence (AI)?

Common definitions

- “The science and engineering of making intelligent machines.” McCarthy (1956)
- “The science of making machines do things that would require intelligence if done by men.” Marvin Minsky (1968) (no relation to Hyman)
- “The study of agents that receive percepts from the environment and perform actions.” Russell and Norvig (2020)
- “AI is the study of how to make computers do things which, at the moment, people do better.” Rich (1983)
- “*Artificial Intelligence is the field of study and technology development that focuses on creating systems capable of performing tasks that typically require human intelligence.*” ChatGPT (2023)

Our notion of AI

- I see it as a *rational maximising agent*
 - Following Russell and Norvig's (2020) taxonomy
- Because that resonates with how economists think
- A computer algorithm performing tasks usually done by humans
- Which differs from machine learning and traditional statistics
- Since it not only provides quantitative analysis
- But also gives recommendations and makes decisions
- While generative AI takes most of the oxygen, other types can be more important in finance

Recommend Russel (2019) for a general overview of AI

Strengths

AI can

- *Interpolate* in highly dimensional spaces
- Identify patterns in large data sets
- Use game theory to become better than humans at dealing with humans
- Be reliable, fast, quick and cheap
- Scale easily
- Reduce human errors
- Provide new insights and solutions
- Enhance security

Weaknesses

- Needs relevant data in its training set
- Usually bad at *extrapolating* — *hallucinates* (later slide)
- May need precise instruction from seemingly unrelated domains
 - Economic theory, history, ethics, politics, psychology, ...
- Humans have markets, firms, and organisations, artefacts that allow them to pool their intelligence to make better decisions. It is not clear that interacting AIs will be able to do so
- We don't know how to incentivise it to align its behaviour with our objectives
- We don't know how AI strategies
- Explainability still in its infancy (later slide)
- For risks, see later Section

Types of AI

1. Reactive machines — basic AI that reacts to specific inputs with predefined responses. Rigid and inflexible, designed for specific tasks
2. Rule-based systems — AI that operates based on predefined rules or logic. Flexible within the domain of rules but not adaptable
3. Limited memory AI (discussed later)
4. Narrow (weak) AI (discussed later)
5. Self-aware AI
6. Artificial general intelligence (AGI)

Our focus is on 3 and 4

5. Self-aware AI — Theory of mind (hypothetical)

- Machines that can interact with humans in an intuitive and human-like manner
- Understand and anticipate human actions based on an understanding of their mental states

6. Artificial general intelligence (AGI) (hypothetical)

- AI capable of understanding, learning, and applying knowledge across a wide range of tasks at a human level
- Extremely complex
- Highly adaptable and generalisable
- AI makes new and better AI — hardware and software
 - Recall Google's AlphaZero
- AGI is an ongoing research goal, with significant advancements anticipated in the coming decades

Is it only a question of data and compute?

Narrow AI vs. Limited Memory AI

- Narrow AI (weak AI)
 - Designed for a specific task or narrow set of tasks
 - Operates based on pre-programmed rules and learned patterns.
 - *Does not update* its parameters with new data after initial training
 - Example: A speech recognition system that doesn't adapt to new accents
- Limited memory AI
 - Can update and adjust parameters based on new data
 - Has a memory that allows for ongoing learning and adaptation
 - More dynamic than narrow AI, often using reinforcement learning or continual learning
 - Example: A self-driving car that improves its performance with more driving experience

Generative AI

- Models that generate new content similar to the data they were trained on
- Applications:
 - Text generation: Chatbots, writing articles and course assignments
 - Image generation: Creating art, design prototypes, deep fakes
 - Music generation: Composing new pieces
 - Data augmentation: Enhancing training datasets for machine learning models, stress testing
 - GPT, DeepArt, DALL-E, Jukedeck, MuseNet, ...

Non-generative AI

- Analyse data, make predictions and classify information without generating new content
 - Predictive analytics, forecasting
 - Natural language processing (NLP): Sentiment analysis, language translation, entity recognition
 - Autonomous systems: Self-driving cars, robotics
 - Image classifiers
 - Speech-to-text systems
 - Recommendation systems
 - Fraud detection systems

Machine learning (ML)

What is Machine learning (ML)?

- Algorithms (models) enabling computers to learn from and make predictions based on data
- Train a model on a dataset so it can do analysis and predictions
- Regressions and other traditional statistical methods are ML
- Useful sources Murphy (2023) and Prince (2023)

Supervised learning

- Model is trained on a labelled dataset
 - $\{P_{AAPL,i}\}_{i=1}^t$
- Each training sample is paired with an output label that the model aims to predict
 - $P_{AAPL,t+1}$
- Model learns to map inputs to outputs by minimising the error between the predicted and the true values
 - $|\hat{P}_{AAPL,t+1} - P_{AAPL,t+1}|$
- Primarily used for classification (predicting categorical labels) and regression (predicting continuous values)
- Common algorithms include linear regression, GARCH, logistic regression, support vector machines, decision trees and k-nearest neighbours

Unsupervised learning

- Model is trained on a dataset without labelled outputs to infer the natural structure of the data
 - Entire text of the *Financial Risk Forecasting* book
- Model tries to learn patterns and structures from the input data without supervision or labelled responses, looking for relationships, groupings and structures in the data
- Used for clustering (grouping data points into clusters) and association (finding rules that describe large portions of the data)

Reinforcement learning


- Involves
 - States — the current situation
 - Actions — choices available
 - Rewards — feedback from the environment
 - Policy — strategy used to determine actions
- Interact with the environment, take action and receive feedback
- Learn to associate actions with feedback and adjust its strategy to maximise cumulative rewards over time
- Used, e.g. in robotics, autonomous driving, game playing (e.g. AlphaGo), resource management, recommendation systems, algorithmic trading, fraud detection, credit scoring and risk management

Neural networks (NN)

- Inspired by the human brain
- Interconnected layers of nodes (neurons) that process data in a hierarchical manner (next two slides)
- Structure
 - The input layer receives the initial data
 - The hidden layer performs computations and extracts features through weighted connections
 - Output layer produces the final prediction or classification
- Learning process
 - Neural networks learn by adjusting the weights of connections based on the error of their predictions
 - Backpropagation, a model's errors are propagated backwards through the layers to update and optimise the weights

Example of a neuron, or node

Receives inputs, applies weights, adds bias, and passes through activation function to produce output




$$y = \sigma \left(\sum_{i=1}^n w_i x_i + b \right)$$

where σ is the activation function, w_i the weights, x_i the inputs and b the bias (constant)

Most common activation function, rectified linear unit (ReLU)

$$\max(0, x)$$

Neural network



Deep learning

- Neural network with many layers
- Each layer transforms the input data into more abstract representations
- Many hidden layers allow the model to learn hierarchical representations of the data
- Automatically extract relevant features from raw data
- Apply non-linear transformations to the data, enabling the network to model complex relationships and patterns
- Require large amounts of data and significant computational resources for training
- Overfitting — controversial, as increasing the size and compute for the same data might or might not have worse out-of-sample performance

Active learning

- Model can query a user or information source to obtain labels for new data points
- Select the most informative data points to label
- Iteratively selects data points, queries for their labels and retrains itself
- By focusing on the most informative examples reduces the amount of labelled data needed, making it useful when labelling data is expensive or time-consuming
- Often involves human experts who provide the labels for the queried instances

Transformer

Vaswani et al. (2023)

- Model weighs the importance of different words in a sequence, regardless of their position
- Injects information (positional encoding) about the position of each word in the sequence, as neural networks can lack an inherent sense of order
- Can handle very large datasets and model complex patterns, making it highly effective for natural language processing tasks
- Applications:
 - Language translation
 - Text generation
 - Sentiment analysis
 - Named entity recognition
 - GPT (Generative Pre-trained Transformer) is a type of transformer

AI Risks and Issues

Main AI societal risks

Weidinger et al. (2022), Bengio et al. (2023), Shevlane et al. (2023), Bostrom (2014), Arthur (1994) and Varian (2018)

1. Trust, embedding and over-reliance
2. Misalignment with human values
3. Malicious use
4. Market power and concentration
5. Bias and discrimination
6. AI feedback loops and self-propagation risks
7. Misinformation and disinformation
8. Autonomy and loss of human control
9. Erosion of privacy
10. Environmental impact
11. Ethical dilemmas in autonomous decision-making

1. Trust, embedding and over-reliance

- AI builds up trust by being good at simple tasks that play to its strength
- As it gets promoted to increasingly sophisticated tasks
- We may end up with the AI version of the Peter principle
- AI presents advice in a way that is hard to reject
 - Ends up a shadow decision-maker
- And becomes essential no matter what senior decision-makers wish
- Usually not credible when someone says
 - “We would never use AI for X”
 - “We will always have a human in the loop”
- So long as it delivers significant cost and efficiency savings

2. Misalignment — Objectives

- How do we ensure AI does what it is supposed to do?
- One of the hardest problems for AI in complex social settings
- Engine needs to be given a precise objective function and cost-benefit calculations
- Impossible to do for problems that are effectively infinitely complex
- Note analogy with the legal system and incomplete contracts
- Object misspecification leads to suboptimal decisions, even bias (3.)
- And disastrous outcomes
 - EURISKO and the canon of human knowledge and ethics
- Meting objective may depend on intuition to find analogous problems from seemingly unrelated domains that can provide creative solutions
 - AI is not good at intuition

1 and 2. Trust and misalignment — Explainability

- We demand human decision-makers explain their decisions
 - Before hiring them, we can ask them to explain how they would react in particular cases
- Explainable AI
 - Transparent and understandable AI decision-making processes
- Essential in many, even most, financial applications
 - Builds trust
 - Accountability
 - Compliance
 - Ethics
 - Model improvement
- Most current AI models have no explainability
- But we are seeing new techniques emerge that help

2. Misalignment — Hallucination

- AI may generate outputs not based on reality
- When asked to extrapolate to cases not in its training dataset
- Because
 - Overgeneralisation, broad assumptions
 - Lack of context, insufficient understanding
- In, for example
 - Text, images
 - Implausible financial stress scenarios that can lead to wrong decisions
 - Beware of GenAI stress testing
- Difficult to detect except for experts
- Undermines trust

2 and 5. Misalignment, bias and discrimination

- Algorithmic bias
 - Erroneous internal representation that biases decision-making away from the desired outcome
- Due to objective misspecification, inadequate training data and/or spurious statistical patterns that change once the AI engine attempts to exploit them
 - Incorrect refusal of products
 - Racial bias
 - Inappropriate language in communication
- Can have serious consequences for the institution employing it
- Key risk holding back the use of generative AI in financial applications
- Significant progress being made

3. Malicious Use of AI

- Use of AI technologies for harmful purposes
 - Find legal or regulatory loopholes
 - Crime
 - Terrorism
 - Nation-state attacks
- Non-financial examples:
 - AI-driven cyberattacks
 - Autonomous weapons systems
 - Manipulation of public opinion or elections
 - Deepfakes for deception or fraud
- Risks to security and stability

4. Market power and concentration

- The technology behind AI
 - Human capital
 - Compute
 - Data
- Currently has increasing returns to scale
- Where all three reinforce the other
- A handful of IT firms end up with the people, data and compute
- Oligopolies and concentration
- Monopolies in specialised areas

AI engines train each other towards undesirable outcomes

- Training data for other engines fed by
 - Visible outcomes of other AIs' advice/decisions
 - Generated data from other AI
- Errors in one AI system propagate to others
- Feeds into
- Feedback loops that reinforce undesirable behaviours
- Can lead
 - Bias amplification
 - Exploitation of weaknesses/loopholes
 - Coordination on undesirable equilibria

AI Economic/Financial System Risks and Issues

Based on
Danielsson and Uthemann (2024)

1. Data

- AI depends on data
- Financial system generates terabytes of data daily
- It is often badly measured and confined to silos
- Little data on most important events as crises are rare (1 in 43 years)
- Observed data lives in the centre of the outcome distribution, not the tails

AI will not learn much about tail events — crises

2. Unknown-unknowns

- All crises have the same code and common fundamentals
 - Liquidity, leverage, one day in 1,000 problem, opacity
- While every crisis is unique in detail
- Axiomatic, as else they (hopefully) would have been prevented
 - We don't know about crises successfully prevented
- Crises are *unknown-unknowns* or uncertain (Knight 1921)
- Little data to train on

AI will not learn much about tail events — crises

3. Endogenous system responses — Setup

- Neural networks are a reduced-form model (“super VAR”) needing policy experiments in their training data (Sims 1980)
- Market participants change their behaviour when regulations/policy changes
 - Lucas (1976) and Goodhart (1974)
- The system changes in response to regulations and crises
- The reaction functions of market participants and the authorities is *hidden* until we encounter stress

AI finds it hard to learn about how the system behaves in crisis as there is little data to train on

3. Endogenous system responses — Wrong way risk

- Neural networks don't learn much from past actions how future stress events play out as it knows little about causal relationships
- Such a problem is almost the opposite of what AI is good for
- When AI is needed most, it knows the least
- AI knowledge is negatively correlated with the importance of decisions

Wrong way risk

4. Objectives — Principal-agent

- Regulations align private incentives with society
- Principal-agent problem between authority and institution (one-sided)
- Becomes two-sided (institution – regulator – AI)
- Regular ways to incentivise — carrots and sticks — don't work with AI
- Need pre-specification of detailed objectives
- Which is not sufficient
- Scheurer, Balesni, and Hobbhahn (2023) insider trading

We don't know how to incentivise AI

4. Objectives — What to tell it to do?

- How to ensure AI follows the objectives desired by its human operators?
- AI needs clear objectives — more than a human because of hallucination and lack of context
- Micro rulebook known and relatively immutable on operational time scales
- Macro rulebook is not
- Mutability increases along with longer time scales and severity
- Practical macro objectives not known
- The worst case is when the objective is unknown ex-ante and cannot be learned, as is the case with the worst financial crises

How to tell AI what to do when we don't know what we want?

4. Objectives — Speed and misalignment

- AI speeds up analysis and decisions
- Particularly problematic in crises
- Human regulators can no longer coordinate rescue efforts and “twist arms”
- See next Section on AI crises
- AI might cause a crisis before its human owner can answer the Fed Chair’s call

How to tell AI what not to do when it acts so fast?

5. Strategic complementarities

- The best course of action for one player increases incentives for others to do the same
 - Multiple equilibria
 - Rapid transition from one equilibrium to another
- Private-sector AI might coordinate on run equilibria
 - Their human owners would prefer to avoid
 - They recognise the damage of a systemic financial crisis to its owner

AI engines train each other towards undesirable outcomes

6. Malicious use

- AI helps those intent on exploiting systemic vulnerabilities
 - By coordinating and manipulating market conditions in a way that benefits them while destabilising the entire market
 - LTCM, Wolkswagen, Hunt brothers, GameStop + Reddit, ...
- Good at pushing into regions of the state space with the strongest complementarities
 - Converge on collusive strategies
 - Jointly chasing the strongest possible bubbles and crashes
 - Find complex cross-market strategies that humans have so far failed to identify
- Rogue traders, criminals, terrorists, nation states

AI helps all the highly resourced agents not concerned about social consequences

7. Risk monoculture

Increasing returns to scale of AI technology

- Risk monoculture drives booms and busts
- Few technology and information companies — oligopolistic, and even monopolistic, market
- Financial data vendors have concentrated considerably over the past few years
- Analytics from the same vendor — procyclicality

It is a concern that neither the competition nor the financial authorities have appreciated the potential for increased systemic risk due to oligopolistic AI technology in the recent wave of data vendor mergers

AI Crises

Based on

Danielsson and Uthemann (2024)

Technology and crises

- One of the earliest applications of the first transatlantic telegram cable in 1858 was the transmission of stock prices
- Nathan Rothschild supposedly used pigeons to get the first news of Napoleon's defeat at Waterloo in 1815 to manipulate the London stock market
- The one-day largest stock market crash in history, on 19 October 1987, was due to algorithmic trading, as was the stress event in June 2007 and several recent flash crashes

How financial institutions optimise

- Maximise profits given the acceptable risk
- Roy's (1952) criterion is useful
- Maximise profits subject to not going bankrupt
- That means financial institutions optimise for profits most of the time, perhaps 999 days out of 1,000
- However, on that one last day, when great upheaval hits the system and a crisis is on the horizon, survival, rather than profit, is what they care most about
- The “one day out of a thousand” problem

Problem

- Options when faced with a shock — run or stay — stabilise or destabilise
- If shock is not too serious, optimal to absorb and even trade against shock
- If avoiding bankruptcy demands a swift, decisive action, such as selling into a falling market, do exactly that

	Run	Stay
Crisis	✓ Right decision	✗ Wrong decision
No crisis	✗ Wrong decision	✓ Right decision

AI to crises

- AI excels at extracting complex patterns from data and, hence, can react faster and with more sophisticated strategies than humans
- AI engines can learn from their competitors
- What this means in practice is that AI engines in private firms and public organisations may end up optimising to influence one another
- Strategic complementarities

If AI decides to stay

- Will not sell or withdraw liquidity
- Buy
- AI is a stabilising force

If it wants to run

- Speed is of the essence
- The first to react gets the best prices
- The last to act faces bankruptcy
- Sell, calls in loans, run others as quickly as possible
- Makes the crisis worse in a vicious cycle
- Extreme volatility
- AI will significantly speed up and strengthen responses, making crises particularly quick and vicious
- AI acts as a crisis amplifier
- *Days or weeks to minutes or hours*

Optimal policy response

Based on

Danielsson and Uthemann (2024)

Relevance

- How the authorities respond will strongly impact the likelihood and severity of crises
- For them to remain relevant overseers of the financial system, they must meet the challenge of
 - Monitoring and regulating private sector AI
 - Address how private AI affect the effectiveness of regulations
 - Face the system risk it poses
 - Harnessing AI for their mission
- If the authorities do not credibly respond, the consequence will likely be more frequent and severe financial crises

Regulation and monitoring today to AI

- Done with
 - PDF reports, database dumps and conversations
- Does not easily allow the authorities to gauge how the private sector would react to
 - Potential regulations
 - Interventions
 - Identify stress resulting from 1 day out of 1,000 problem
 - Particularly from institution interactions
- AI can open up a new dimension in the authority-private sector interaction
 - Making regulations more robust and efficient
 - Allows better monitoring of misconduct and stress
 - Help in finding latent systemic risk
 - Design more efficient crisis interventions

Options

- A.** Implement their own AI engines
- B.** Set up AI-to-AI links
- C.** Triggered central bank standing facilities
- D.** Public-private partnerships

A. Authority AI system — Setup

- Set up their own AI engines
 - To design regulations
 - Evaluate the effectiveness of interventions
 - Gain expertise in AI
- Key characteristics
 - Speed
 - Objective specification

A. Authority AI system — Issues

- Such a system would not withstand the Lucas critique
- It would hence struggle with tail events
- However if
 - Well-designed
 - Informed both by
 - Confidential data
 - AI-to-AI benchmarking links discussed below,
- Would be of considerable benefit to both the micro and macro authorities

B. AI-to-AI links — Setup

- AI-to-AI links
- Authority AI directly talks to AI in private firms and other authorities
 - Legal framework
 - API
- Benefits
 - Regulating private-sector AI
 - Benchmarking
 - Test compliance
 - Perhaps how they decide on a loan application
 - Respond to a shock
 - And a crisis intervention
- Might not need data sharing

B AI-to-AI links — Gaming

- How would the system react
 - Lucas (1976) and Goodhart (1974)
- Depends on the effectiveness of the authority's setup
 - How it asks the private sector how it might respond to planned policy reactions
- To counteract gaming
 - Same policy intervention with a number of private AI
 - To ascertain whether a particular institution is attempting to game
 - Send a number of alternative queries
 - More difficult for the private AI to strategise

B AI-to-AI links — Hurdles

- No technological hurdles
- The data and sovereignty issues are more difficult
- But AI-to-AI API links focus on reactions and not data
- Benchmarking might not need access to the underlying technology
- Responses might be sufficient

C. Triggered central bank standing facilities

- Central banks prefer discretionary facilities — *Constructive ambiguity*
- Too slow in AI crises (recall the last section)
- Facilities with predetermined rules for immediate triggered responses
- Can rule out bad equilibria
 - If private agents know CBs will intervene if prices drop by X
 - Will not coordinate on strategies that are only profitable if prices drop $> X$
 - Keeps moral hazard low
- If poorly designed, allow gaming and, hence, moral hazard

D. Public-private partnerships — Setup

- Authorities need AI that matches private-sector AI
- Hard to do in-house
 - Conservative, political and bureaucratic culture
 - Employment policies particularly difficult
 - High salaries of AI experts
 - Junior staff with high authority
- Outsourcing more likely
 - Private firms set up and run engines and AI-to-AI links
 - As already is increasingly common, e.g. in fraud, KYC, AML, credit risk, etc.

D. Public-private partnerships — Downsides

- Oligopolistic, even monopolistic, AI market structure
- Prevent the authorities from collecting information
 - Decision-making processes
- Private sector firms also prefer to keep technology proprietary and not disclose it, even to the authorities
- Lack of control
- Lack of knowledge of AI
 - But, does not prevent them from regulating derivatives without much derivative expertise

Applications

Regulatory design and supervision

- Design
 - Query private firms on hypothetical customer interactions
- Supervision
 - Query responses to routine interactions

Monitor stress

- Ask for stress scenarios
- Simulate market stress scenarios
- Aggregate responses and feedback across industry
- However, stress scenarios will likely continue to be backwards-looking, hence not improving the problem of the lack of tail observations

Interventions

- Simulations of alternative interventions via iterations
 - Iteratively query private sector AI and those of other agencies
 - Aggregate and feed responses back
- Find the best approach for
 - Liquidity injections
 - Forbearance
 - Bailouts
- Identify the impact on market stability, credit availability, and overall economic health
- Allows for more informed decisions
- The mere presence of such an arrangement might act as a stabilising force in a crisis if perceived as competent and credible

Planning

- All four options need to be in place before the next stress event
- Compel the private sector to comply — private sector will resist
 - *Might* see benefits in cooperating to prevent/resolve crises
- Stronger links only be activated in an emergency

Kill switches

- Now suspend trading
- Not as viable for AI
- System might not function effectively if a key AI engine is turned off
 - Significant market disruptions
- Difficult to predict AI responses to shutdown commands
 - might not account for the interconnected nature of modern financial markets where multiple AI systems interact
 - AI systems are designed to learn and adapt to their environment.
- Difficult to design a kill switch that can be effectively deployed without additional harm

Bibliography I

Arthur, W. Brian. 1994. *Increasing Returns and Path Dependence in the Economy*. Ann Arbor, MI: University of Michigan Press. ISBN: 978-0472064960.

Bengio, Yoshua, Geoffrey Hinton, Andrew Yao, Dawn Song, Pieter Abbeel, Yuval Noah Harari, Ya-Qin Zhang, Lan Xue, Shai Shalev-Shwartz, Gillian Hadfield, et al. 2023. "Managing AI risks in an era of rapid progress." *arXiv preprint arXiv:2310.17688*.

Bostrom, Nick. 2014. *Superintelligence: Paths, Dangers, Strategies*. Oxford, UK: Oxford University Press. ISBN: 978-0199678112.

ChatGPT. 2023. *ChatGPT's Definition of Artificial Intelligence*. <https://www.openai.com/chatgpt>.

Danielsson, Jon, and Andreas Uthemann. 2024. *Artificial intelligence and financial crises*.

Goodhart, Charles A. E. 1974. *Public lecture at the Reserve bank of Australia*.

Knight, Frank. 1921. *Risk, Uncertainty and Profit*. Houghton Mifflin.

Lucas, Robert E. 1976. "Econometric Policy Evaluation: A Critique." *Journal of Monetary Economics* 1.2 Supplementary:19–46.

McCarthy, John. 1956. "Dartmouth Conference: The Birth of Artificial Intelligence." In *Proceedings of the Dartmouth Conference*.

