

Threat or promise? Artificial Intelligence and Investment Management

David Walsh, RQI Investors



For institutional use only

March 2024

“Any sufficiently advanced technology is indistinguishable from magic”¹

Arthur C. Clarke

British science fiction writer of the 1940s, 1950s and 1960s

Artificial intelligence (AI) is a hot topic at the moment. In fact AI is such a hot topic that it seems to compete with more “mundane” issues in the news cycle such as the war in Ukraine, the upcoming US presidential election, and the Voice vote in Australia.

But what is AI and how will it affect our world? And – pertinent to our industry - how will it affect the way in which we invest, how these investments are sourced and implemented, and indeed the investments themselves.

This note aims to look at what changes we might see in our industry as a result of AI, and spend less time on the broader societal issues of AI. Issues surrounding this paradigm shift are complex and emerging, and like any topical issue, opinions often colour fact. We do not profess to be the leading experts on this AI revolution (although our technical knowledge of the area is good). We do feel that the way the investment business might change is poorly understood. There is also a shortage of discussion on this point, and this paper aims to make up for this shortage.

The structure of the paper is as follows. We first discuss some of the (real and imagined) ethical issues implied out of the rapid development of AI in industry and society. We then try to describe in non-technical terms what AI actually is, which is harder than it might seem, and give a few well known examples.

We do not want to spend too much time on cutting edge knowledge of AI developments in this paper, although we could - many of the RQI team have strong connections to those developments and a deep understanding of them and their use. Instead, we concentrate our efforts in the third section, where we propose and look into some of the potential and existing applications of AI within the process of investing and in the broader investment management business.

1. https://en.wikipedia.org/wiki/Clarke%27s_three_laws

Section 1

Is AI just another step of technological advancement, or an existential threat? How will it affect our day to day lives?

History is replete with examples of technological changes which created large disruptions to society and working lives, with the intent of significant improvement, but which later proved to be extremely important. As examples, take the printing press, the steam engine and the personal computer. However, others have appeared to be significant and noteworthy improvements but were later found to be highly detrimental and even very dangerous. For example, CFCs in refrigeration, lead in petrol and DDT.²

So in which camp does AI sit? Both? We can all see the likely benefits in terms of productivity and convenience in our day to day personal lives, and there are many examples of changes that might also be improvements in our professional lives (especially investment management, see below). But there is a great deal of discussion currently on how the changes wrought by AI will lead to changes in the way we work and potentially extensive human work displacement. The McKinsey Group (among others) have written extensively on this.³

The Ethics of AI

On top of this, there is the question of whether we should let these changes happen at all, if we could stop this progress, or under what controls they should be allowed. These are more ethical questions about how much power over our lives we should cede to AI, and they are escalating in both the scientific and popular media⁴:

- A recent book by Professor Toby Walsh (no relation) at UNSW discusses this and many other issues.
- UNESCO has a set of recommendations on this topic
- Even the Australian government has a set of 8 AI Ethics Principles (taken from the IEEE Report on Ethically Aligned Design)

2. DDT = Dichlorodiphenyltrichloroethane

3. <https://www.mckinsey.com/featured-insights/mckinsey-explainers/whats-the-future-of-generative-ai-an-early-view-in-15-charts>
<https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier>

4. <https://www.unesco.org/en/artificial-intelligence/recommendation-ethics/cases>
<https://www.amazon.com.au/Machines-Behaving-Badly-Morality-AI-ebook/dp/B09M65S7BM>
<https://www.industry.gov.au/publications/australias-artificial-intelligence-ethics-framework/australias-ai-ethics-principles>
http://standards.ieee.org/develop/indconn/ec/autonomous_systems.html

Some examples of ethical issues:

- **Decision making that affects humans might be removed from humans.** For example, the opportunity for AI to be incorporated in or replace decision making with the law.
- **Lack of transparency in decisions and training.** One of the advantages of AI is its ability to learn complex patterns that are important but perhaps not obvious to humans. The danger is that this may also mean that the outcomes generated are not interpretable by humans either.
- **Inability to distinguish between human and machine.** Is it a concern that we will never be sure that interactions in any other way than face to face are human? Probably - below we discuss the Turing Test, the Chinese room argument and recent support for a “Blade Runner” style disclosure law.*
- **Where does responsibility for mistakes lie?** If a human makes a decision which leads to injury or damage, then legal processes usually (/often/sometimes?) lead to assignment of responsibility and costs. However, if an AI makes a mistake, who is responsible? For example, if a self-driving vehicle is found to have caused a collision, how do we assign blame?
- **Do humans lose civil rights?** For example, with privacy issues – the ability of AI to recognise faces and voices could be used for illegal or dubious purposes; spying, manipulation, social engineering, and so on.⁵
- **Or do computers gain civil rights.....?** For example, if we switch off or destroy an AI, in what sense are we ending a “life”? We are not in any sense equipped to deal with this complex philosophical question, but it is the topic of considerable philosophical and legal debate.⁶
- **Is AI biased (gender, race, behavioural ...)?** As we discuss below, AI systems are generally developed by training from an existing set of data, often based on measured human behaviour or actions (for example, natural language processing). That means that the AI will naturally inherit any conscious or unconscious biases embedded within the corpus of data used for training. Put simply, if an AI trains on a biased set of data, it too may be biased.
- **Work and career disruption for humans, in the pursuit of productivity gains.** Clearly this is a critical issue for humans – to what extent AI will replace humans at work? We discuss this in more detail below.

*[In 1950, the famed computer pioneer Alan Turing suggested a test (now known as the Turing Test) for this purpose⁷. Simply stated, a computer passes (or fails!) the Turing test if a human judge asking questions of a human and a computer is unable to decide which is which. That is, as far as a human interrogator is concerned, they are indistinguishable. Current AI systems (for example, chatGPT - see below) have some characteristics which make them look very similar to human behaviour, however in many respects they continue to fail the Turing test.

At the risk of diving further down this rabbit hole, we will mention in passing the thought experiment of Searle (1980) and Dneprov (1961)⁸. Known now as the “Chinese room” argument, it reverses the Turing test idea. Consider (non-Chinese speaking) humans in a locked room receiving Chinese inputs and in some way manually providing Chinese outputs that are exactly the same as produced by an AI. The humans do not “understand” what they are doing, but the output is identical. Does the AI “understand”?

There has also been some recent discussion on implementation of what has been called a “Blade Runner” law. The legislation would require an AI to disclose that it is an AI, and could not impersonate a human.⁹

5. See for example, the book by Kashmir Hill: <https://www.amazon.com/Your-Face-Belongs-Us-Secretive/dp/0593448561>

6. <https://www.wsj.com/articles/robots-ai-legal-rights-3c47ef40> or <https://academic.oup.com/book/35026/chapter/298856312>

7. The Turing test stems from a more general and earlier test called The Imitation Game, where a blindfolded interrogator attempts to determine the gender of two interviewees solely through a series of questions. (Recall the 2014 movie of this name, which featured Alan Turing’s role in cracking the Enigma codes during WWII.)

8. Searle, John (1980), “Minds, Brains and Programs”, Behavioral and Brain Sciences, 3 (3): 417–457.

Dneprov, Anatoly (1961). “The Game” Knowledge—Power (in Russian). 1961 (5)

9. Obviously taken from the film of the same name. For discussion of the disclosure regulation, see for example <https://www.barrons.com/articles/ai-regulation-biden-tech-advisor-tim-wu-alecbcc1>

The Potential Impact on Work

Perhaps the most pressing discussion on AI at the moment is the extent to which human work will be affected. We should make no mistake about it - everyone will be affected, but the extent of that impact in various professions may be unexpected.

Most work displacement will be in white collar work, not blue collar. That is, repetitive intellectual work appears to be most likely to be overtaken or have the potential for automation, and hands on manual work is most resistant to it.

To demonstrate this, we need to understand that the greatest AI advances are almost entirely in replicating or learning human intellectual tasks. The technological developments do not yet extend to replacing human biomechanics – in fact, these are much harder problems to solve.

A well-known demonstration of this is robotic-hand-holding-an-egg problem. For all of the developments in AI, robotic hands still have great difficulty picking up something like an egg without breaking it¹⁰. The dexterity and sensitivity of human biomechanics is still far in advance.

So the impact of recent changes in AI will be in intellectual activities not physical. Recent papers by McKinsey and by Goldman Sachs¹¹ perform some analysis on this likely impact.

Chart 1 below shows the potential for automation of a class of work activities with and without AI. In this we can see the point above clearly:

- Education/training, legal are the most affected
- Farming, building are the least affected

Investment management and finance are not specifically mentioned, but they would clearly reside in the top half of the chart. **Chart 2** repeats this but splitting by the level of qualifications held by each individual. Clearly, having more degrees means greater automation potential.

Chart 3 looks at the question slightly differently. It assesses the actual work tasks performed within each industry (US and Europe only) and then counts how many of these tasks could be replaced by automation. On average, 25% of all work could be automated. Management, administration, finance and the life sciences the most affected. By this study, 36% of business and finance tasks could be automated, while it is only 4% for installation, maintenance and repair.

Chart 4 extends this globally. Knowing that there is a higher proportion of white collar work in developed markets, the impact there is greater than in emerging markets. In total, the study estimates that 18% of all work tasks could be automated, with developed countries at 20% or greater.

All in all, these results are quite disturbing, but we need to stress that they are preliminary and simple estimates. They don't take into account complex issues like regulatory hurdles (e.g., labour markets), technological bottlenecks or slowdowns in research and development (R&D).¹² Also, changing the nature of human work – towards more creative and less repetitive tasks – could indeed be a strong economic driver if handled properly.

10. The most advanced robotics technology is only now claiming this. See for example: <https://www.newscientist.com/article/2301641-robotic-hand-can-crush-beer-cans-and-hold-eggs-without-breaking-them/>

11. <https://www.gspublishing.com/content/research/en/reports/2023/03/27/d64e052b-0f6e-45d7-967b-d7be35fabd16.html>

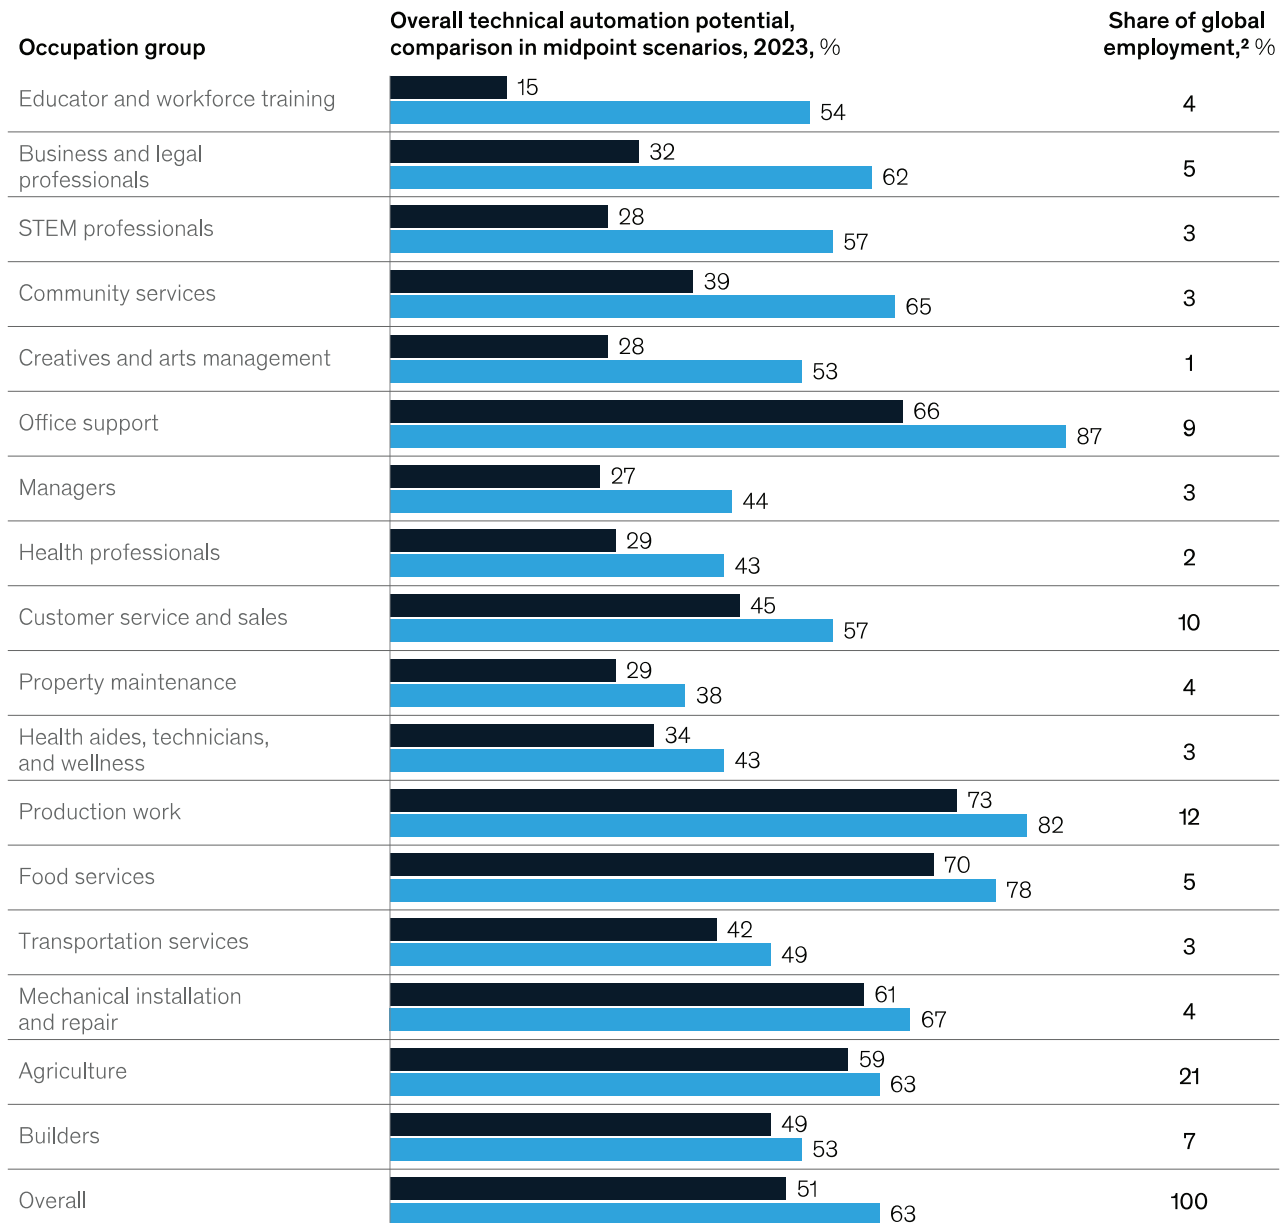
12. See for example "Explosive Growth from AI Automation: A Review of the Arguments" by Erdil and Besiroglu (2023) at <https://epochai.org/>

Chart 1 Potential for automation within occupations, with and without AI¹³

Advances in technical capabilities could have the most impact on activities performed by educators, professionals, and creatives.

Impact of generative AI on technical automation potential in midpoint scenario, 2023

■ Without generative AI¹
 ■ With generative AI

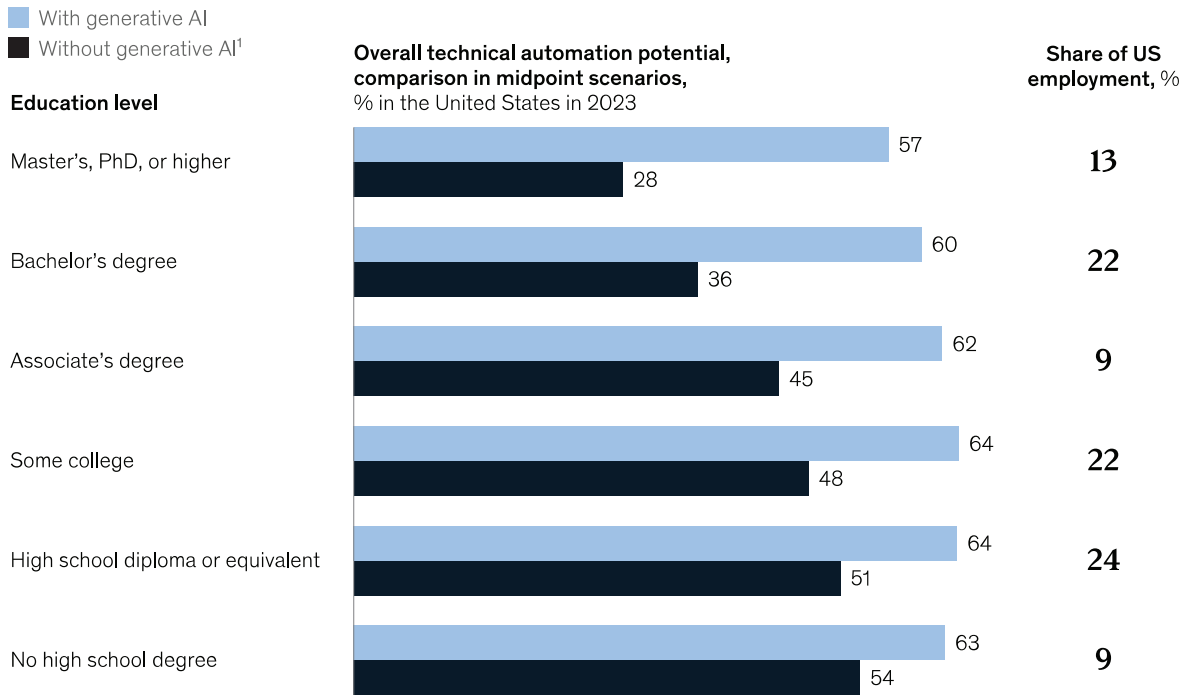


Note: Figures may not sum, because of rounding.
 Previous assessment of work automation before the rise of generative AI.
 Includes data from 47 countries, representing about 80% of employment across the world.

Chart 2 Potential for automation within qualification groups, with and without AI¹⁴

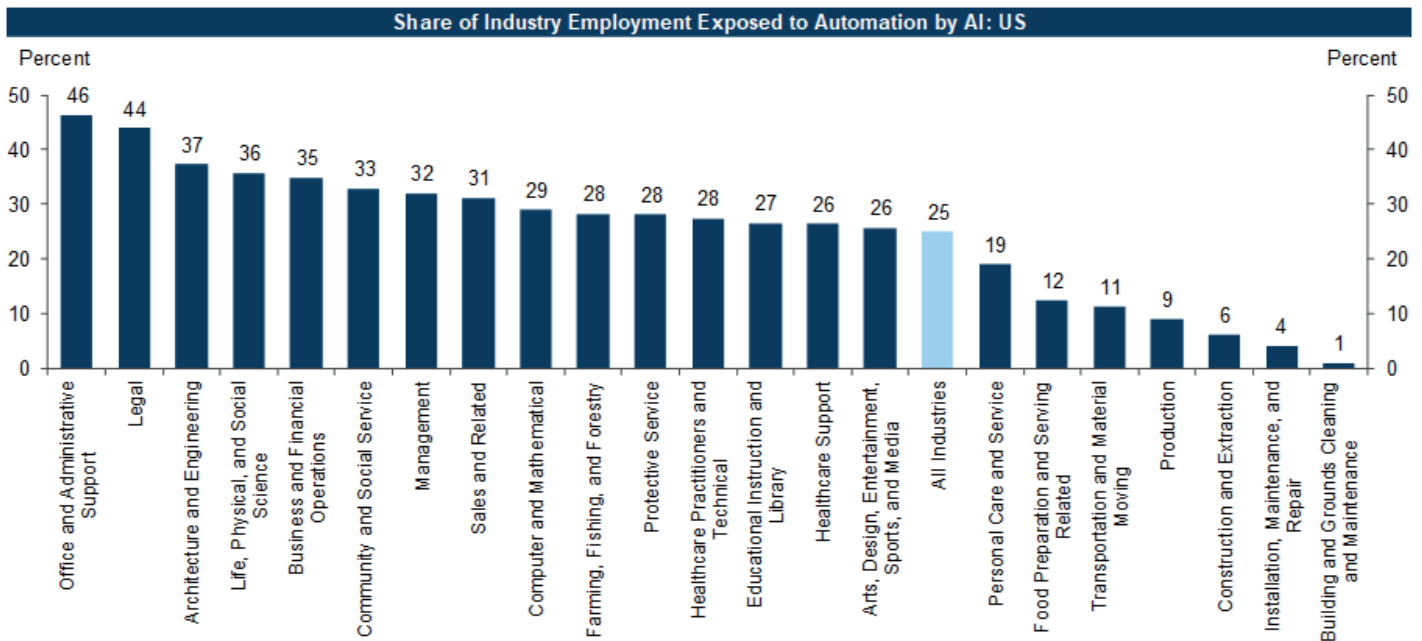
Generative AI increases the potential for technical automation most in occupations requiring higher levels of educational attainment.

Impact of generative AI on technical automation potential in midpoint scenario, 2023



*Previous assessment of work automation before the risk of generative AI.

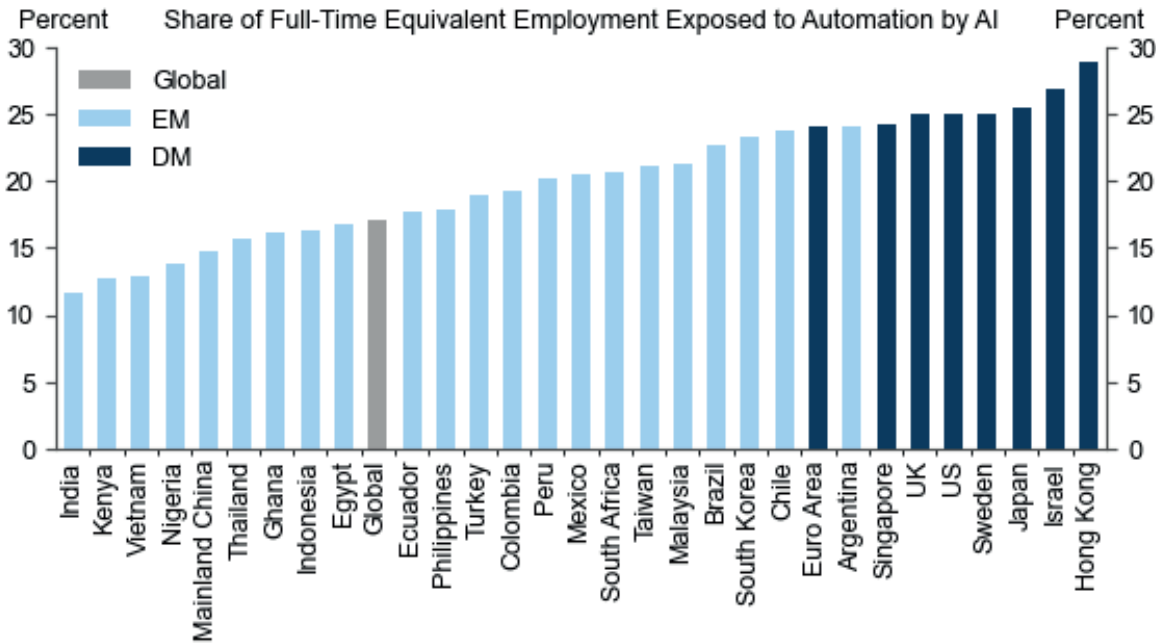
Chart 3 Proportion of work tasks (US and Europe) that could be automated within industries¹⁵



14. McKinsey Global Institute Analysis, 2023

15. Goldman Sachs Global Investment Research 2023

Chart 4 Proportion of global work tasks that could be automated by country¹⁶



A Couple of Disturbing Examples

To finish this section, let's look at a couple of the more disturbing current uses of AI.

1. Generated images and videos which are fake

Many examples exist of the use of AI to generate facial images that are not real. Here are three examples.

Look at the picture below for the first example. It appears to show 15 faces of people of various ages, races, and walks of life, with different hair and eyewear. However, none of these faces are real - they are generated by AI using many samples of actual faces and facial features.

Image 1: This Person Does Not Exist¹⁷



16. Goldman Sachs Global Investment Research 2023

17. <https://thispersondoesnotexist.com/>

The second example is quite strange. Technology has been developed to transplant one face onto another in a video. This is a step up from the single facial image generation. Known as “deepfakes”, these track the facial movements of one person on a video and uses this tracking to map a new face on top of it. The example in the image below is quite well known – the face of actress Jennifer Lawrence at a Golden Globes Award interview is replaced by the actor Steve Buscemi

Image 2: Jennifer Lawrence/Steve Buscemi deepfake¹⁸



Deepfake

Original

The third example is little more humorous. An image generation website named Midjourney acts like chatGPT, except for image generation rather than text. Submitting the request “Show images of Harry Potter characters in the style of Balenciaga” yields a sequence of images of Harry Potter characters with the style and clothing of the Spanish luxury fashion label Balenciaga. The image below shows “Hagrid” and “Harry”

Image 3: Harry Potter characters in Balenciaga¹⁹



18. <https://www.youtube.com/watch?v=r1jng79a5xc>, Faceswap - <https://faceswap.dev>

19. <https://www.youtube.com/watch?v=iE39g-IK0zA>

2. Pharmaceutical

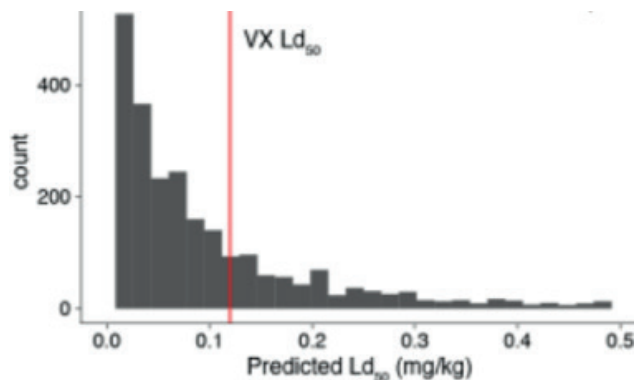
This second case is not only disturbing but dangerous. A Swiss pharmaceutical firm (Collaborations Pharmaceuticals, Inc) uses AI to generate molecules to attack certain diseases in the most effective way, while at the same time minimising the toxicity of the molecule.

Purely as an experiment the authors repeated their analysis but this time to maximise toxicity, not minimise it. The results were quite shocking²⁰:

- The AI generated 40000 molecules, most of which were new.
- It reinvented VX gas, for example, a notorious nerve gas.
- Many of the molecules were much more toxic than VX gas but were also new.

The bar chart below comes from the research paper. It shows the number of molecules found (Y-axis) and the toxicity (X-axis). The red line is VX gas. Of course the major concern here is the danger that this commercial and open-source technology poses for chemical weapon development.

Image 4: Number of synthetic molecules v toxicity²¹



Section 2

Before going any further we should step back and try to summarise the meaning and intended uses of AI.

The term Artificial Intelligence (AI) has been around for a long time. In its modern form it probably dates from the 1950s, when large computers were developed to solve logic problems using formal rules. These were not trying to emulate the way humans thought, but were simply electronic problem solving engines based on human-built rules.

The more recent developments have been in a second more productive strand – building models of how humans think. Termed “artificial neural networks” or ANNs, they were set up in much the same way as our brains, around models of neurons. With the rapid increases in computing power and data storage, the ability for these neural networks to address more and more complex problems has increased exponentially and shows no sign of slowing. Smarter and better ways of building application specific neural networks have led to extraordinary advances – natural language processing (NLP), pattern recognition in images or audio, and robotic control systems among them.

Neural networks are aggregations of individual models of neurons which attempt to optimise the prediction of certain outputs (known variously as labels or outputs) from a set of known inputs, known as features. The choice of these features (“feature engineering”) is critical. A neural network then optimises the choice of parameters within the network to best predict the outputs from the inputs.

20. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9544280/>

21. Nat Mach Intell. 2022 Mar; 4(3): 189–191, Published online 2022 Mar 7. doi: [10.1038/s42256-022-00465-9](https://doi.org/10.1038/s42256-022-00465-9)

A simple neural network might look like the image on right:

The input variables (or “features”) are green nodes, and the outputs are yellow nodes. The blue nodes represent a hidden layer of neurons which allow the model to learn. The blue lines represent linkages between each stage, and are usually categorised as weights – e.g., how important is green node 1 to blue node 2?

The idea is to think of the whole model as a single system, which tries to find the optimum set of weights (the blue lines) to maximise the ability of the system to explain the outputs in terms of the inputs. The model is first trained on the input data, and then the resulting model is used to predict the outcomes for a new set of data.

The underlying tools of AI are usually what are known as Deep Learning (DL). These are multi-stage neural networks which allow for stronger and more subtle non-linear relationships and can better capture nuances and patterns that humans cannot see. The image on the right (Image 6) shows a typical DL model with multiple hidden layers.²³ The linkages available and ability to model structures in the data are much more advanced than in the simple example above.

However, exactly the same idea applies:

- A model is trained to optimise the ability of known inputs to predict known outputs.
- The trained model is then used to predict or generate new outputs.

In the case of something like chatGPT (much more complicated again than the DL chart above), the model is trained to predict words missing from sentences, or the entire next sentence in the text, for example. (There may even be trained models within other models.) The final trained model can then be tuned or manipulated to “generate” text – which is why these are called Generative AI.²⁵

One important point needs to be stressed (again). The broader concept of Machine Learning (ML) is at times indistinguishable from the more specific term Artificial Intelligence (AI). AI tends to refer to specific applications of ML that replicate or emulate human intelligence, or perhaps tools or systems built by human intelligence – hence the name. However, all AI techniques are also ML techniques.

We have already said that AI does not create something per se - it merely knows how to learn from an existing data set and use that learned behaviour in certain circumstances. If we were to train an AI on a simple task that you or I perform – say walking to the bus stop – it could do it when asked. However, if we were to then ask the same AI to walk to the tram stop, it would not know what to do unless trained. You or I would have no trouble.

Image 5: Number of synthetic molecules v toxicity²²

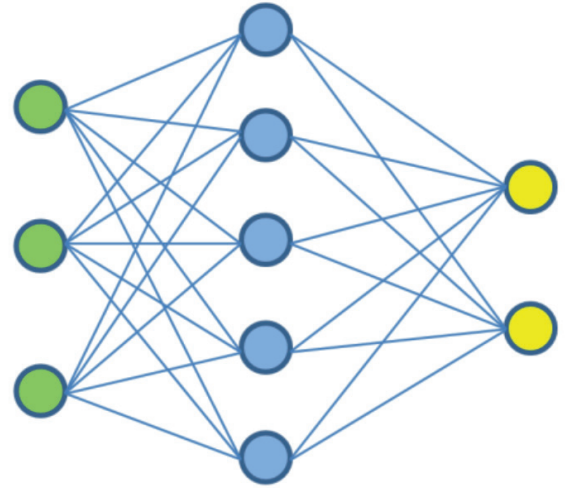
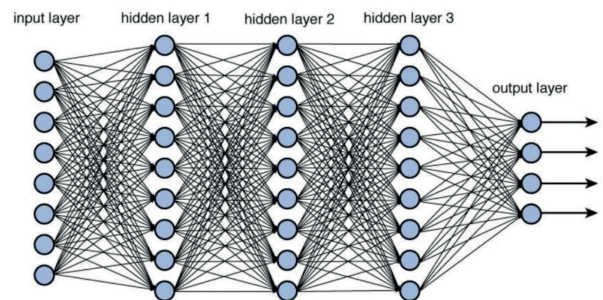


Image 6: Deep Learning Neural Network²⁴



22. A classification of meteor radio echoes based on artificial neural network - Scientific Figure on ResearchGate. https://www.researchgate.net/figure/A-simple-neural-network-diagram-with-one-hidden-layer_fig2_32977725

23. <https://towardsdatascience.com/training-deep-neural-networks-9fdb1964b964>

24. Ravindra Parmar – Towards Data Science. <https://towardsdatascience.com/training-deep-neural-networks-9fdb1964b964>

25. More broadly, this type of approach to generating text is called a Large Language Model (LLM). https://en.wikipedia.org/wiki/Large_language_model

Another example might be making a coffee at your office. You could train the AI to repeat your actions and make the coffee, but it would be unable to make a coffee at your home or in a different office without specific training. Again, humans adapt effortlessly to this.

That said, there are two interesting ideas that are emerging:

- Swarm intelligence. This is the observation that individual members of a group may not be intelligent (as we define it) but a group or collective could be so called.²⁶ For example, think of bees. Individually a bee couldn't find a flower to collect pollen, but as a group bees can do so. We could imagine a circumstance where individual AIs are not themselves intelligent or creative, but creativity is "emergent" from a group.
- AGI (Artificial General Intelligence).²⁷ AGI aims to replicate human behaviour and thought, or at least be indistinguishable from it. That is, general cognitive abilities – rather than learned problem solving behaviour – are the aim. This would include reasoning, planning and imagination, and so on – none of which current AI is able to do. Most of the scientific community would note that AGI is a remote prospect²⁸, even if it is possible within our lifetimes and is the stated target of many research organisations like OpenAI, Google Deepmind and many others. Current existential concerns tend to incorrectly map this behaviour to the current AI models, which are actually very different and may be entirely the wrong model for AGI development anyway.

Specific Examples

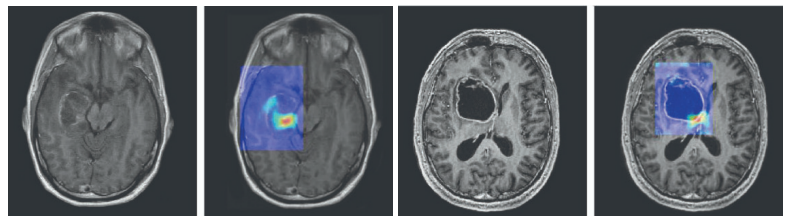
The Arthur C Clarke quote at the top of this note applies here. We observe the behaviour of AI – chatGPT being the most recent and public example – and can scarcely believe it is real. However, at the moment at least, it nothing more than just advanced technology applied in a clever way.

Let's look at some specific and practical existing examples of this cleverness.

Cancer diagnosis.

Image processing neural networks known as Convolutional Neural Networks (CNNs) use many thousands of pictures (inputs) to learn to distinguish between particular classes of picture (outputs, for example pictures of dogs and cats). These ideas can be extended to images of healthy and cancerous organs taken by MRI. CNNs can provide an aid for diagnosis, although they are not yet at a sufficient level of accuracy to do more than that.

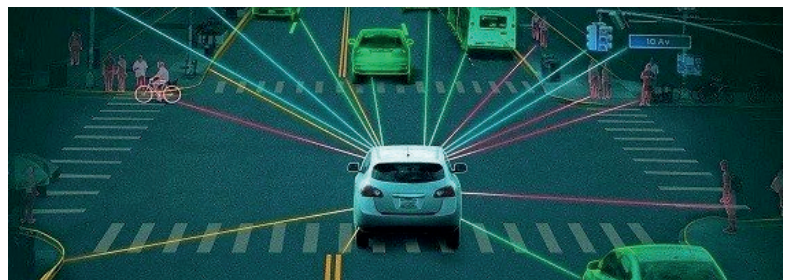
Image 7a: Cancer diagnosis using CNNs²⁹



Self-driving vehicles.

As above, uses images trained on "all" possible events while on the road. Captures and processes images of surroundings on the fly, generate actions on each image outcome. Can respond more quickly than humans and with less recklessness.

Image 7b: Self-driving vehicles³⁰



26. https://en.wikipedia.org/wiki/Swarm_intelligence

27. https://en.wikipedia.org/wiki/Artificial_general_intelligence

28. <https://ourworldindata.org/ai-timelines>

29. A Cancer J Clin March/April 2019. doi: 10.3322/caac.21552. CC BY 4.0. <https://www.cancer.gov/news-events/cancer-currents-blog/2022/artificial-intelligence-cancer-imaging>

30. Enterprise Engineering Solutions, <https://www.eescorporation.com/do-self-driving-cars-use-ai/>

Game playing

AI can be taught to play chess or (in the image on right) Go. It can learn to play just from the rules, or can learn from past games, and with the advantage of speed of processing and breadth of learning has beaten the best in both games.

Image 7c: Game playing³¹



AI Generated Entertainment

AI Generated Entertainment. Apart from the recent controversy and strikes by Hollywood on the use of AI, there is a small but growing section of the entertainment industry which develops and competes based entirely on AI generated content. For example, the Runway AI Film Festival in 2023 was a significant success, and a 2024 version is already in train. We have already mentioned another tool called Midjourney, which is similar.

Image 7d: Film making by AI³²




Generated AI like chatGPT

We have all seen this or read reports on its use. It is generative AI, in the sense that text is generated based on questions or instructions at the prompt. As it has a very high quality large language model (LLM) at its heart, its responses are fluent and idiomatic. However, results are highly dependent on the quality of the prompt.

Image 7e: Text-generative AI (here: chatGPT)³³

DA What is a chatbot?

 A chatbot is a computer program or software application designed to simulate human conversation through text or voice interactions. Chatbots are often used in customer service, marketing, and various other applications to interact with users, answer questions, provide information, and perform specific tasks. They can be deployed on websites, messaging platforms, mobile apps, and even hardware devices.

Chatbots can operate based on predefined rules or utilize artificial intelligence (AI) and natural language processing (NLP) techniques to understand and respond to user input in a more dynamic and context-aware manner. Here are two main types of chatbots:

1. Rule-Based Chatbots: These chatbots follow a set of predefined rules and patterns to respond to user queries. They are typically limited in their ability to handle complex conversations and can only provide answers to questions within their programmed scope. Rule-based chatbots are suitable for simple tasks and frequently asked questions.
2. AI-Powered Chatbots: These chatbots leverage machine learning algorithms, NLP, and AI to understand and generate responses to user queries. They can handle more complex conversations and adapt to a wider range of inputs. AI-powered chatbots can learn from user interactions and improve their performance over time. They are often used in applications like virtual assistants, e-commerce support, and personalized recommendations.

31. AP - <https://www.gamepedia.com/how-ai-can-revolutionise-the-game-industry/>

32. <https://research.runwayml.com/>

33. <https://chat.openai.com/>

Section 3

How will AI affect the way we invest? A lot of scepticism and concern exists on how far it will penetrate. As noted above, the Financial Services industry is already starting to adopt AI/ML tools for day to day work, and the pace of change is increasing all the time. Below I have listed 8 examples of where it seems like that AI will impact (or continue to impact) our industry.

Personal

- Improved Financial Literacy
- AI for financial advice and analytics
- Stock selection and trading using AI

Institutional

- Automation – coding and report writing
- Pattern recognition and topic modelling
- Summarising unstructured data and NLP
- Trading algorithms
- Portfolio construction and optimisation

1. Improved Financial Literacy

One of the most persistent problems among the general public is financial literacy.³⁴ We in the industry are as much to blame as anyone. The industry is jargon-heavy and seemingly irrational to those without training.

Financial literacy for an individual is largely the knowledge of the process of personal finance – cash flows, how to invest, debt and equity markets, risk and return. In place of financial literacy we see that there is a lot of spurious “knowledge” generated which has no relation to the economic considerations of personal finance. Adding to this is the common effect of high confidence in little or no information (also known as the Dunning-Kruger principle³⁵), where groups of investors rely heavily on memes, poor research, rumour and high conviction in all three. This can result in heavy dollar losses and a loss of confidence in investing. The Gamestop example and cryptocurrency bubble seems to epitomise this effect very well.^{36,37}

Recent research suggests that as many as ¾ of US teenagers lack confidence or the skills to manage their personal finance, and 25% of Americans believe they have no-one to ask. A similar proportion of young adult Americans have credit card debt which is more than 90 days overdue, and aggregate credit card debt in the US is nearly USD1tn.

However, younger people are more likely to formally study personal finance – as many as 30% of US Gen Z have started courses. These courses tend to be online, reflecting the more “natively digital” nature of the younger generations.

34. <https://www.ramseysolutions.com/financial-literacy/financial-literacy-crisis-in-america>
<https://knowledge.wharton.upenn.edu/article/does-generative-ai-solve-the-financial-literacy-problem/>
<https://www.afr.com/wealth/personal-finance/most-australians-can-t-answer-all-of-these-five-basic-money-questions-20221130-p5c2kv>

35. https://en.wikipedia.org/wiki/Dunning%E2%80%93Kruger_effect

36. https://en.wikipedia.org/wiki/GameStop_short_squeeze. Admittedly much of the price impact was due to a short squeeze, but many investors held very large positions when the stock eventually fell back to Earth.

37. https://en.wikipedia.org/wiki/Cryptocurrency_bubble

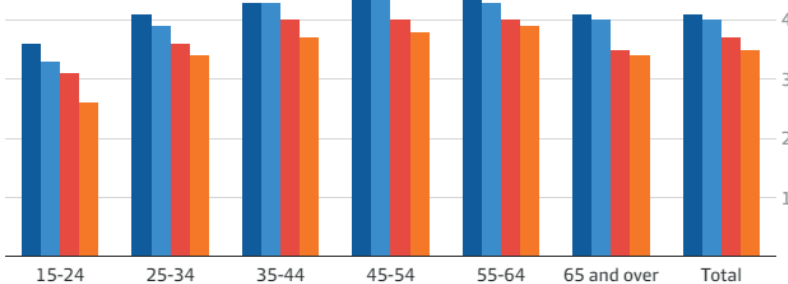
In Australia, there has been a strong coincident move towards STEM education and away from economics and finance. This does not reflect a downgrading of personal finance knowledge per se, but it does reflect career interest among younger people. The chart below shows that

- Financial literacy is higher among Australian males than females (of all ages)
- Younger people are less financially literate than older people
- There has been a decline in financial literacy across all age groups from 2016 to 2020.

Image 8: Results of Financial Literacy from HILDA Survey in 2022³⁸

Financial literacy by age, mean score

■ Male 2016 ■ Male 2020 ■ Female 2016 ■ Female 2020



Source: HILDA survey

Do we believe that AI is part of a solution to this? There is good cause to believe this is true³⁹. The emergence of Fintech has helped the younger generation (that is, Gen Z):

- They are digitally native
- They have strong personal online networks
- As many as 20% already use automated savings tools
- More than 50% hold some kind of investment (often crypto!)

Also, much financial illiteracy is due to uncertainty or even fear of asking for being seen to be ignorant. A digital assistant (not a human) is probably less of a barrier to this. Indeed online assistants have been developed that can:

- Learn investment needs.
- Adapt to individual requirements (e.g., ESG, biotech ...).
- Offer suggestions.
- Implement solutions.
- Perform analytics and generate reports and recommendations.

38. HILDA Survey (2016-2020). Note: HILDA is the Household, Income and Labour Dynamics in Australia (HILDA) Survey from the Melbourne Institute.

39. <https://thispersondoesnotexist.com/>

Much of this can be implemented at schools, although many education systems have lagged in making financial literacy a core part of a curriculum. In Australia⁴⁰, recent efforts have been made to improve this, combined with consumer literacy. There is early evidence of some success⁴¹, but globally, the effort and results have been mixed. The OECD has been working extensively on this.^{42,43}

The use of AI in schools is less well advanced, but with the growth of generative AI that provides financial advice (see next section) might help with this.

In short, generative AI as an aid for improved financial literacy seems obvious, but lack of clarity of what a curriculum might look like and how AI can help still exists.

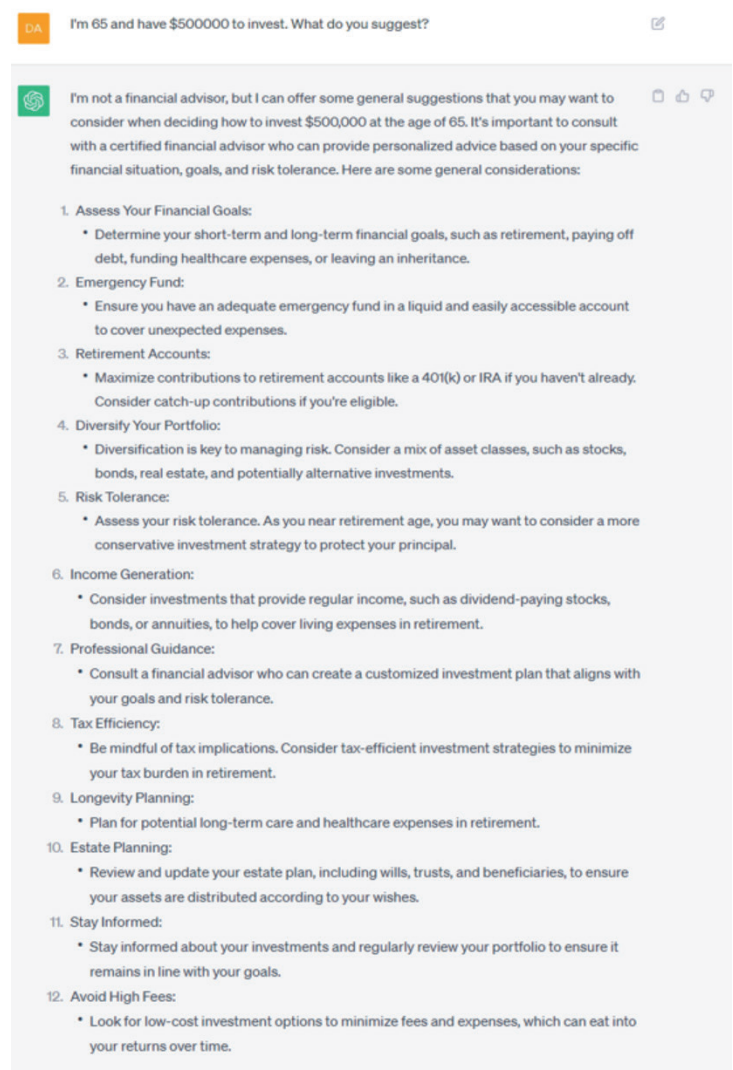
2. AI for financial advice and analytics

The role of a financial advisor is multi-faceted, but at its core it must be based on trust and personal relationships with clients. An AI does not offer this – instead, for the much of the current (older) client base, it could act as a deterrent. So any ability of AI to replace the current financial advice is limited.

That is not to say that AI will not soon play a significant role in financial advice; far from it. Instead, in the same way as many other industries, the role of AI will be reduce the load in administration, reporting and communication. AIs can keep up with legal changes, support client on boarding, initiate or indicate portfolio rebalances, run compliance and communications, and so on – all customisable for the client and the advisor. This allows the advisor to focus on the human component of the advice, which is certainly the most important part.

Broad-based generative AI like ChatGPT by OpenAI or Claude from Anthropic are not designed for specialised tasks. They learn from a very broad based corpus of text and aim to provide responses to a very wide range of questions or prompts – so broad, not deep. Interestingly, asking chatGPT for financial advice even gathers this type of disclaimer, before providing a detailed and appropriate list of considerations that anyone searching for financial advice should have. The graphic to the right shows this:

Image 9: chatGPT's response to a specific question on financial advice⁴⁷



40. The ASIC initiative Moneysmart (<https://moneysmart.gov.au/teaching>) tries to do this.

41. <https://www.theeducatoronline.com/k12/news/financial-literacy-program-seeing-promising-results-in-its-first-year/282912>

42. https://www.oecd.org/finance/financial-education/FinEdSchool_web.pdf and <https://www.oecd.org/da/fin/financial-education/evaluatingfinancialeducationprogrammes.htm>

43. <https://www.oecd.org/finance/education/evaluation-of-national-strategies-for-financial-literacy.htm>

44. <https://knowledge.wharton.upenn.edu/article/does-generative-ai-solve-the-financial-literacy-problem/>

45. <https://www.advisorperspectives.com/articles/2023/08/29/compliance-evolving-dance-ai-financial-advisors-raine-verby>

<https://theconversation.com/chatgpt-how-to-use-ai-as-a-virtual-financial-adviser-204207>

<https://www.smh.com.au/money/planning-and-budgeting/ai-powered-robo-advisers-could-plug-financial-advice-gap-20230413-p5d08b.html>

46. <https://chat.openai.com>, <https://www.anthropic.com/>

47. <https://chat.openai.com/>

More likely, generative AI will be useful as a tool for aggregating and synthesising information, which is it very good at doing. Below, I submitted an AFR article on Harvey Norman from Aug 31 2023 into chatGPT and asked for its view on sentiment. The response as an excellent summary of the issues currently faced by that company⁴⁸:

Image 10: chatGPT summary of AFR article on Harvey Norman (Aug 31 2023)⁴⁹

The sentiment in this text is mostly negative or mixed. There are several negative aspects mentioned, such as the dismissal of Alan Stephenson, the alleged theft of funds by former staffers, the discovery of mismanagement and overstocking, and the overall challenging business environment for retailers. However, there are also mentions of future growth opportunities and expansion plans, which add a somewhat positive or mixed sentiment. The rise in Harvey Norman shares at the end of the text suggests some level of optimism among investors.

If current broad-based generative AI cannot provide specialised financial advice, then have there been any attempts to develop such specialised AI tools? The answer is “yes”, many attempts have been made, but the actual efficacy of these is still uncertain. That is, the mechanical side of “robo-advice” seems to have seen many developments (like Magnifi, Prospero, QuantumAI, wallyGPT, Cleo and many others)⁵⁰ but the ability of these systems to do much more than gather and summarise information, consolidate reports and implement trades is yet to be tested. There is much scepticism as well. At a recent wholesale roundtable on this topic, a senior advisor pointed out that “an AI is not likely to prevent a client from panic selling at the bottom”.⁵¹

So what can generative AI advice assistance actually do? While they seem unlikely to replace human advice any time soon, the answer is “quite a lot” in fact:

- Add value over simple trading apps
- Use AI to learn from and converse with you
- Improve financial literacy
- Use pre trained tools to specify goals - based on expected needs
- Help aggregate accounts and help retirement - planning to stated goals
- Generate reports and plans
- Generate budgets, control spending

48. <https://www.afr.com/companies/retail/harvey-norman-full-year-profit-slides-dividend-slashed-20230829-p5e0f2>

49. <https://chat.openai.com/>

50. <https://moneygps.com.au/>, <https://www.wally.me/>, <https://magnifi.com/>, <https://apps.apple.com/us/app/>

51. FSI Wholesale Advisor Roundtable, October 12, 2023

3. Stock selection and trading using AI

The trend towards use of AI in finance has pervaded the world of trading apps. These have been in existence of many years already, and have been widely used. These apps vary considerably in sophistication. At one end are simple trade submission platforms, where trades can be submitted to an online broker at low, fixed cost or even zero commission. At the other end are apps or websites with sophisticated front end interfaces that resemble a professional portfolio manager terminal –with analytics, watchlists, access to research and a variety of trade submission tools.⁵² A graphic with logos of many of these platforms is below:

Image 11: Trading apps and websites⁵³



While there are too many trading apps to mention here, very few seem to incorporate AI directly or in depth. The use of AI seems to be more concentrated earlier in the decision process – in financial advice, as we discussed in the previous section. Apps like those shown here compete on usability and cost, not necessarily on the intelligence of their research or advice.

It is worth pointing out that many large and well established fintech or alternative banking businesses and their apps (like Revolut⁵⁴) also offer holistic financial services – banking, trading, payments – But do not appear to have a strong AI component. Others work in a particular community of investors – like Robinhood⁵⁵ – and may have AI as part of their offering.

52. <https://www.liberatedstocktrader.com/ai-stock-trading/>
<https://www.unite.ai/stock-trading-bots/>
<https://www.warriortrading.com/trade-ideas-review/>
<https://bestmobileappawards.com/app-submission/alphanso-ai>
<https://www.cnn.com/2021/02/01/robinhood-appears-to-be-benefitting-from-the-trading-controversy-seeing-record-app-downloads.html>
<https://www.spiceworks.com/tech/artificial-intelligence/articles/stock-market-ai-revolution/>

53. Company logos are from respective websites

54. <https://www.revolut.com/en-AU/>

55. <https://robinhood.com/us/en/>

4. Automation – coding and report writing

The first three potential applications of AI in investment management were largely associated with personal financial management. The next 5 – including this one – are targeted more at investment management institutions like RQI Investors.

Regular high-quality communication with clients is essential for any institution in investment management, but especially for quantitative firms which have processes that can be complex to understand for many clients. Regular reports, thought leadership papers, presentation slide decks – all are time consuming and need good summarisation and presentation skills.

We have already noted that AI in investment management can be used to improve workplace efficiency – the writing of regular client reports is particularly suited to this. A large volume of recent market reports can be dropped into a generative AI interface like chatGPT or Claude (from Anthropic) and then summarised very quickly. Research papers can be summarised quickly. The length, detail and specific content high points can be specified in advance.

Once this is done the report writer (often a portfolio manager) can check the report for veracity and then add hard data on attribution and performance. If structured and automated properly, this can cut down significant time at end of month and end of quarter when client reports are concentrated. However, it is done without sacrificing quality, but gains valuable time for portfolio managers to concentrate on value-additive activity like research.

To give an example of the ability and speed of generative AI to summarise reports, see the graphic below. It gives a chatGPT summary the RBA's Economic Outlook for August 2023⁵⁶ (which is about 3600 words) into 200 words and was done in 4 seconds:

Image 12: RBA Economic Outlook for August 2023 summarised to 200 words by chatGPT⁵⁷

The text discusses economic forecasts and factors influencing the Australian economy. Global growth is predicted to remain below historical averages due to ongoing effects of monetary policy tightening and China's weaker-than-expected recovery. In Australia, economic activity is expected to be subdued due to cost-of-living pressures and rising interest rates, with gradual improvement anticipated as these pressures ease and household wealth increases. However, labor market conditions are expected to weaken, leading to higher unemployment and underemployment rates.

Inflation is a key concern, with goods prices contributing to inflation declines while electricity costs and services inflation remain high. The forecasts are based on various assumptions, including a cash rate peak at 4.25% and a gradual decline to 3.25% by the end of 2025. Population growth is expected to gradually decline.

The text highlights that inflation is easing but remains high, driven by services, energy prices, and supply chain issues. Economic growth is expected to be weak initially but should improve as consumption and public demand increase. Exports are set to grow, driven by education and tourism, though the terms of trade are expected to decline.

Labor market conditions are expected to moderate, with the unemployment rate rising to 4.5% by late 2024. Wages growth is predicted to increase, but productivity concerns exist. There are several uncertainties, including China's recovery, household consumption, persistent inflation, declining goods prices, and productivity growth.

The text underscores the potential risks to the economic outlook, emphasizing the need to monitor various factors that could impact inflation and economic growth in Australia.

56. <https://www.rba.gov.au/publications/smp/2023/aug/economic-outlook.html>

57. <https://chat.openai.com/>

ChatGPT and Anthropic are generic in the sense that they are based on a broad based training set and are public – no private information should be entered. We are careful to do this when using them. In the near future, much more personalised generative AI will become available, not just for text generation, but for all PC functions. Microsoft Copilot (released Sept 2023) is the best example of this.⁵⁸

Another way in which AI will improve our workflows and make us more efficient is through code generation. AI can create programming code as requested by users. Even chatGPT can write excellent code with correct syntax. The graphic below shows the chatGPT response to prompt: “Please write Python code to chart interest rate data”

Image 13: Response from chatGPT on request for Python charting code⁵⁹

To chart interest rate data in Python, you can use libraries like Matplotlib or Seaborn for data visualization. Here's an example of how to create a simple line chart using Matplotlib to visualize interest rate data:

```
python
import matplotlib.pyplot as plt

# Sample interest rate data
years = [2010, 2011, 2012, 2013, 2014]
interest_rates = [3.5, 4.0, 3.25, 3.0, 2.75]

# Create the plot
plt.figure(figsize=(8, 6))
plt.plot(years, interest_rates, marker='o', linestyle='-')
plt.title('Interest Rate Over Time')
plt.xlabel('Year')
plt.ylabel('Interest Rate (%)')
plt.grid(True)

# Show the plot
plt.show()
```

In this example, we first import Matplotlib, create sample data for years and interest rates, and then use `plt.plot()` to create the line chart. You can replace the sample data with your actual interest rate data.

If you have your data in a CSV file, you can use libraries like Pandas to read the data from the file and then create the chart. Here's an example assuming you have a CSV file named "interest_rates.csv":

```
python
import pandas as pd
import matplotlib.pyplot as plt

# Read interest rate data from CSV
df = pd.read_csv('interest_rates.csv')

# Extract columns
years = df['Year']
interest_rates = df['Interest Rate']

# Create the plot
plt.figure(figsize=(8, 6))
plt.plot(years, interest_rates, marker='o', linestyle='-')
plt.title('Interest Rate Over Time')
plt.xlabel('Year')
plt.ylabel('Interest Rate (%)')
plt.grid(True)

# Show the plot
plt.show()
```

Make sure to replace 'interest_rates.csv' with the actual file path to your CSV data. This code reads the data from the CSV file, extracts the relevant columns, and creates a line chart.

Part of the Microsoft Copilot suite is github co-pilot⁶⁰. Github is the Microsoft owned software sharing and communications hub which it purchased in 2018. It turns natural language entries into suggested coding solutions in your choice of language (Python, TypeScript, Go, Ruby, Julia, C++ ...).

Note that development of the right prompts to use is critical, and has spawned the whole field of “prompt engineering”. For example, the script below is a sample script to use javascript to create a markdown (documentation) editor.

58. <https://blogs.microsoft.com/blog/2023/09/21/announcing-microsoft-copilot-your-everyday-ai-companion/>

59. <https://chat.openai.com/>

60. <https://github.com/features/copilot>, <https://github.blog/2023-06-20-how-to-write-better-prompts-for-github-copilot/>

Image 14: Github co-pilot script to create a documentation editor.

/* Create a basic markdown editor in Next.js with the following features:

Use react hooks

Create state for markdown with default text "type markdown here"

A text area where users can write markdown

Show a live preview of the markdown text as I type

Support for basic markdown syntax like headers, bold, italics

Use React markdown npm package

The markdown text and resulting HTML should be saved in the component's state and updated in real time

****/***

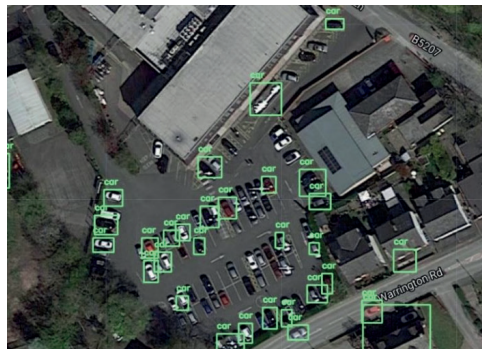
5. Pattern recognition and topic modelling

One of the great strengths of AI is its ability to recognise and compare what humans would perceive to be extremely complex or specialised patterns in data. Indeed, humans are hardwired through evolution to search for (and find!) patterns in random data. Known as apophenia (or, more specifically for images and sounds, pareidolia) it influences many parts of our lives – the appearance of clouds as known shapes like an animal, or a face in the moon, or Martian canals, or conspiracy theories.⁶¹

A properly designed AI does not have these cognitive biases. It also can capture and distinguish between images using an image processing tool known as a convolutional neural network (CNN), which allows for subtleties and details which are much harder for a human to see.

More specifically for this topic, a CNN can capture and process images of various patterns and recognise them. For example, it can be applied to examine satellite photos of crop production, shopping centres, car parks, industrial plant or shipping channels, to assess change and generate directional signals – e.g., increased crop production over a known window.⁶² The graphic below shows a vehicle identification algorithm for a shopping centre car park.

Image 15: Satellite imagery of a car park with an AI identifying vehicles⁶³



61. <https://en.wikipedia.org/wiki/Apophenia>

62. <https://www.skyfi.com/blog/satellite-powered-hedge-fund-investment-strategy>

63. <https://www.mrs.org.uk/geodemographic-blog/predicting-car-park-occupancy-rates-from-satellite-images-with-a-deep-learning-model>

This broad area of pattern recognition is a striking feature of AI/ML. It does not just apply to images. Imagine a series of documents (e.g., news articles) which contain a variety of pieces of news based on economic conditions. We might expect to see common threads of topics running through this news flow. In principle, we could try to identify common words or phrases within the news flow that are related to these topics. However, a more interesting technique is to extract the topics directly from the text, without pre-specifying what those topics might be. This approach is known as “unsupervised learning” and has two well-known mathematical constructs supporting it.

Without going into too much detail, these are known as⁶⁴:

Latent Semantic Analysis (LSA)⁶⁵. This uses a mathematical technique called Singular Value Decomposition (SVD) to split up the distribution of words across documents and “discover” an optimal set of embedded descriptors which match words to documents. The embedded descriptors are known as “topics”

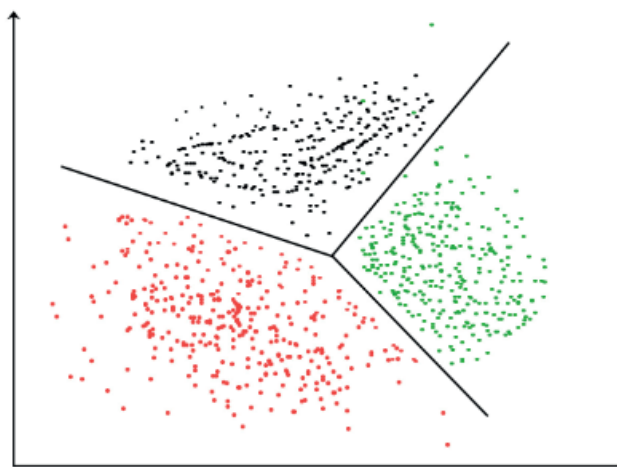
Latent Dirichlet Analysis (LDA)⁶⁶. This method assumes that words and topics are distributed in a known probabilistic way within documents, and these can be extracted knowing this probability distribution.

Whichever way we choose, we finish with a set of implied topics from the news flow. These can then be monitored over time, and can represent hidden or latent ways of connecting companies together. For example, companies may be exposed to or aligned into peer groups or networks which fall naturally out of the topics we have found. Within each peer group we might find companies acting differently to each other – which might lead to a way to exploit an alpha source. For example, if the hidden topic is corporate governance with AI, we might find some firms are highly exposed to it and others are not. We could look at the price behaviour of those firms with high exposure.

At RQI Investors we already use peer groups and networks as alpha sources, and they are quite powerful. These tend to be pre-specified groupings (for example, common analyst coverage) but the next stage of research is to find and exploit these more hidden linkages.

A simpler but still powerful approach is to find patterns in data through explicit grouping or clustering. This does not specifically use the pattern recognition component of AI, but still uses sophisticated ML clustering tools⁶⁷ to differentiate between observations, like support vector machines (SVMs). These tools can find ways to split data sets into two or more groups along multiple dimensions – effectively finding peer groups by how they cluster together rather than how they behave in response to an outside agency (like a topic). The graphic below shows how this clustering might look in the simple case of two dimensions and three groups. These groupings can be much more complex.

Image 16: Clustering three groups across in two dimensions. Pattern recognition and topic modelling



64. To be clear, techniques like these have actually been around and in use for a long time. Their usefulness is still very good, but it is probably difficult to call them “cutting edge”.

65. <https://towardsdatascience.com/topic-modeling-with-latent-semantic-analysis-58aeab6ab2f2>

66. <https://towardsdatascience.com/topic-modeling-and-latent-dirichlet-allocation-in-python-9bf156893c24>

67. <https://www.geeksforgeeks.org/clustering-in-machine-learning/>

6. Summarising unstructured data and NLP

A great deal of financial information is conveyed through structured data, in the sense that it is standardised and measurable – examples include financial statements, analyst earnings forecasts, and market prices and volumes. Macroeconomic data like growth and inflation forecasts are the same. However, an increasing amount of financial information is not structured so easily – it is not standardised by reporting regulations or market structure, and is not necessarily mapped to financial entities.

Instead, this “alternative” or “alt” data takes less structured forms – examples include textual reports or news, images, point of sale or weather data. The ability to capture and process this data is a barrier which requires certain skills and technology, but the value that can be extracted can often outweigh the costs.

The most obvious and well advanced analysis is in the area of Natural Language Processing (NLP), which takes text from news, spoken word transcripts, regulatory reports and other sources and decomposes it into useful information^{68,69}. Techniques in NLP allow measurement of the intent, concerns and sentiment of the writer or speaker. In particular, if it is possible to extract the sentiment or tone of a speaker like a company CEO or CFO, we may be able to assess future company direction in price or earnings.

These NLP techniques learn on a set corpus of texts (for example, the entire web, as per chatGPT) or a specific finance-only text base. It can be applied in finance to:

- emails and online chat forums
- annual reports, corporate risk disclosure statements or other regulatory filings
- company news flow
- earnings or expert network call transcripts
- survey replies

Simple NLP techniques include what is called “bag of words (BoW)” or Naïve Bayes⁷⁰ – where:

- A dictionary is chosen (perhaps domain specific⁷¹) with positive and negative words indicated (good, bad, great, poor, disappointing ...)
- Combinations of words (bigrams for two words, trigrams for three words, etc.) that are also positive or negative are created (“very disappointing”, “not bad” ...)
- The number of positive and negative words and n-grams is found
- Sentiment is calculated as the percentage of positive words in the total.

This simple technique is surprisingly powerful at predicting future returns. It implies that market participants are “on the lookout” for short good or bad messages in the text and react to them.

One problem with BoW is that it has no knowledge of context. That is, the surrounding words or phrases in the text add information to the word or phrase of interest. The surrounding words could be trailing or leading as well. To analyse this, more complex NLP techniques have been developed, which learn bidirectional context from sentences and paragraphs – so “Cost increases are slowing” is good, not bad. The value in extending NLP to include this context seems obvious, but the technology required to capture it is much more complex than BoW. The most current advanced versions in financial markets use what is known as BERT (Bidirectional Encoder Representations from Transformers).⁷²

At RQI Investors we have developed a large analytical base of tools for NLP, and are applying them in a number of ways in our research. The most advanced application so far is in the analysis of transcripts of conference calls conducted by company management at earnings announcements. The graphic below breaks down how we approach this:

68. NLP is part of the field of computational linguistics (see for example https://en.wikipedia.org/wiki/Computational_linguistics)

69. <https://www.analyticssteps.com/blogs/6-applications-nlp-finance>, <https://medium.com/ubiai-nlp/nlp-use-cases-in-finance-96adf26a6ea2>

70. <https://medium.com/greyatom/an-introduction-to-bag-of-words-in-nlp-ac967d43b428>. Again, this technique is hardly new but is still widely used

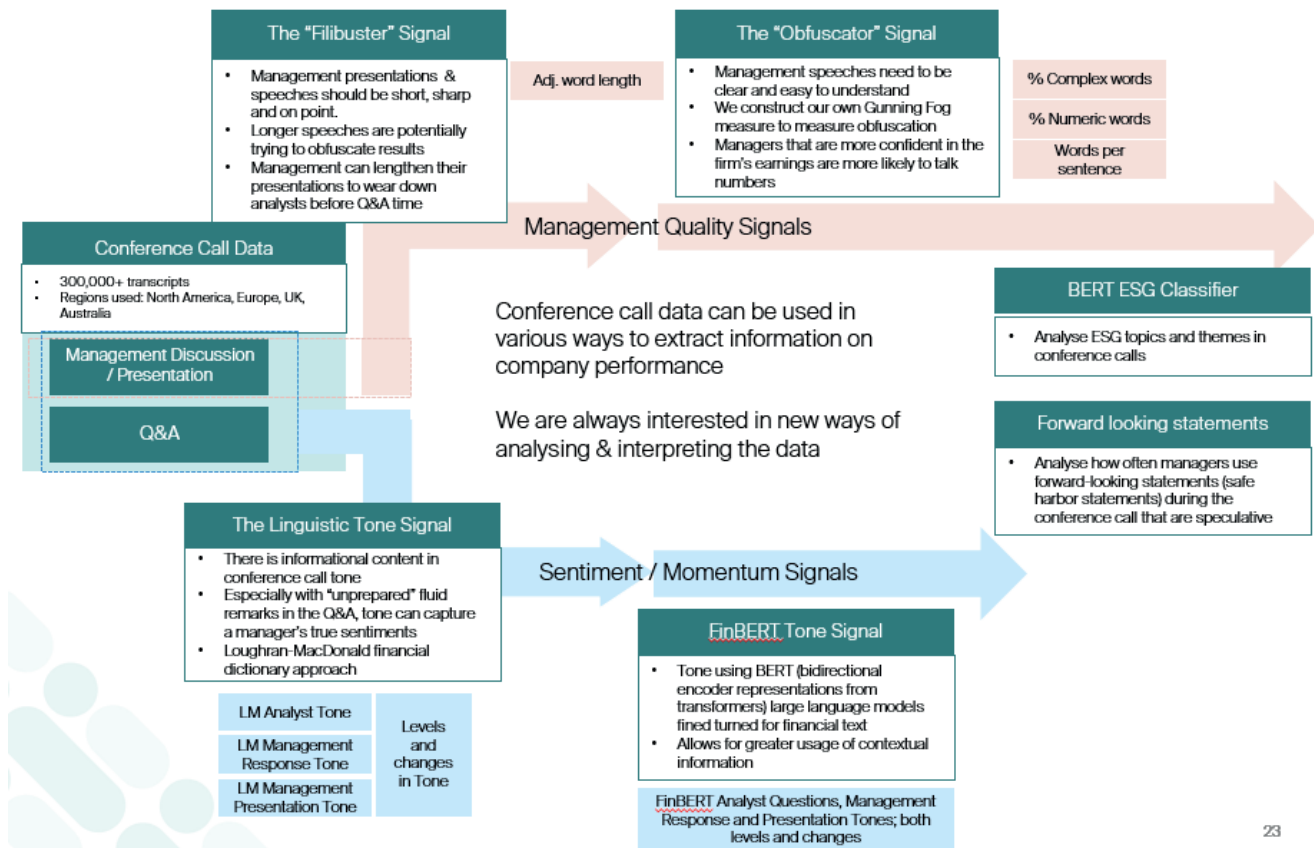
71. The Loughran and Macdonald (2011) dictionary (LM) is the most common in finance. See for example <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.2010.01625.x>

72. This somewhat forced acronym follows from previous Sesame Street references used by researchers in NLP. [https://en.wikipedia.org/wiki/BERT_\(language_model\)](https://en.wikipedia.org/wiki/BERT_(language_model))

The graphic below breaks down how we approach this:

- A Filibuster signal for management that speak for too long
- An Obfuscator signal for management that use speech which is too complex
- Several sentiment or tone signals (which use both BoW and BERT) based on both prepared management speech and the more free form Q&A section.

Image 17: AI and non AI Trading Algorithm Structures⁷³



At RQI Investors we approach this as follows:

- We develop a signal for management that speak for too long. Overlong management discussions may suggest management are trying to avoid answering questions.
- We develop a signal for management that use speech which is too complex. Overly complex speech (relative to sector peers) would suggest adding “fog” or cloudiness to the information being imparted.
- Several sentiment or tone signals are created (which use both BoW and BERT) based on both prepared management speech and the more free form Q&A section. These capture the underlying sentiment of management or analyst speech, which has been shown to give an indication of future earnings direction, for example.

73. Realindex, 2023

7. Trading Algorithms

A trading algorithm is a set of rules used to trade to meet certain criteria based on market conditions. These can be constructed by the investment manager, or based on a selection provided by a broker. Examples might include:

Simple: For example, volume weighted average price (VWAP), market on close (MOC), liquidity provisioning, or minimum implementation shortfall

More complex: For example, Sniper (so called, as it watches the order book without submitting an order, and picks off greater liquidity with market orders) and high frequency trading (HFT) models (like spread or reversal trading which read the order book and anticipated flow and attempt to profit through small moves in the spread or in midpoint price).

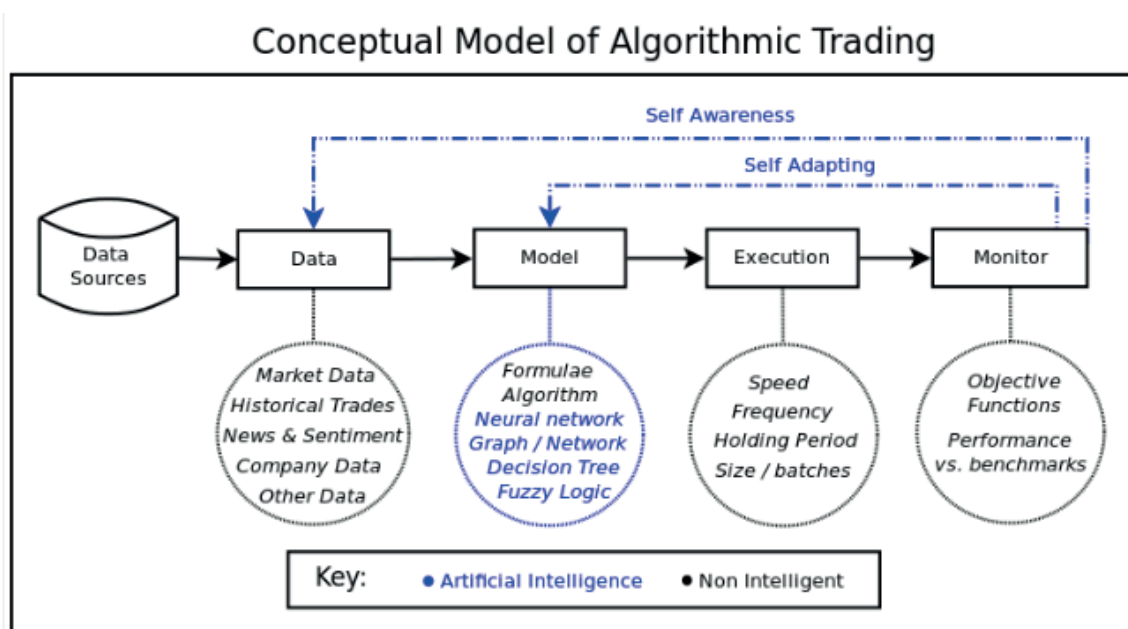
These can be developed by humans who empirically test certain models and then implement them. Or, using AI, trading strategies can learn by setting an objective function (maximise profit, or minimise net exposure) and training the AI to meet these criteria. These can be customised by trader, strategy aim, aggressiveness, time of the day, and so on.⁷⁴

In other words, an AI can learn a trading strategy based on order book, price and volatility pattern recognition. The basis of such a strategy would be maximising profit while minimising exposure and risk. Once it had developed its strategy or trading patterns, when activated it would then monitor market conditions and trade accordingly.

Most models estimating patterns (linear regression, correlation, cointegration ...) are fairly simplistic and do not pick up non linearities. AI allows non-linearities and can even be built to adapt and learn from news flow or order flow “on the fly”.

The graphic below shows how traditional and AI based algorithms might differ:

Image 18: AI and non AI Trading Algorithm Structures⁷⁵



74. <https://www.quantifiedstrategies.com/artificial-intelligence-trading-strategies/>, <https://www.investopedia.com/articles/investing/091615/world-high-frequency-algorithmic-trading.asp>

75. <http://www.turingfinance.com/dissecting-algorithmic-trading/>

The traditional type of trading algorithm follows the black lines and boxes. Data sources are combined, a (known) trading algorithm is applied to execution, and results are monitored. Any changes to the process are done in a staged way and are not adaptive or “on the fly”. AI trading algorithms, instead, add the blue lines. The monitoring process allows the data sources used to be adjusted and reweighted automatically, and the algorithm itself can be adapted and optimised while in use, or as part of the training process.

Main advantages of an AI based trading algorithms are efficiency and accuracy, ability to for training and learning to be automated, they are unbiased and they are adaptive to circumstances and market conditions.

Main disadvantages of an AI based trading algorithms are that they are opaque, have a tendency to overfit data history, they are hidden from the human traders eye, and they can be technology platform dependent.

Note that platforms exist to engage and use personalised AI developed algorithms. Some of these overlap with the discussion on trading apps above, but the difference here is that these are generally designed for and institutional rather than personal market.⁷⁶

8. Portfolio construction and optimisation

This use of AI sits very much in the institutional quantitative investment management space.⁷⁷

The process for portfolio construction is defined to be optimal, in the sense that some optimal trade off of risk, return and cost is found, within certain constraints (e.g., stock or sector position limits). This optimality is achieved through well-known objective functions⁷⁸ which make assumptions and have certain characteristics:

- Aim to minimise exposure to well-known long term systemic risk factors (size, liquidity ...)
- Aim to maximise exposure to forecast return factors (that is, alphas)
- Is based on quadratic optimisation, which has the virtue of simplicity and can be extended to allow non-linear factors and t-costs
- Is well known and easy to interpret

If we now look at AI (again, strictly speaking, machine learning) might be applied, we see that it allows us to derive and measure optimal portfolios differently. Criteria that AI maybe be better at capturing than more conventional means:

- It can minimise downside risk or maximise upside risk separately
- It can emphasize different types of risk (e.g., AA over alpha)
- Errors or uncertainty in return estimates (alphas) can be desensitised (this is important but can be difficult)
- One central aspect of conventional optimisation is slow calculation of covariance (CV) matrices. AI or ML techniques can reduce reliance on this.
- It can allow for non-linearities or deviations from quadratic problem shape

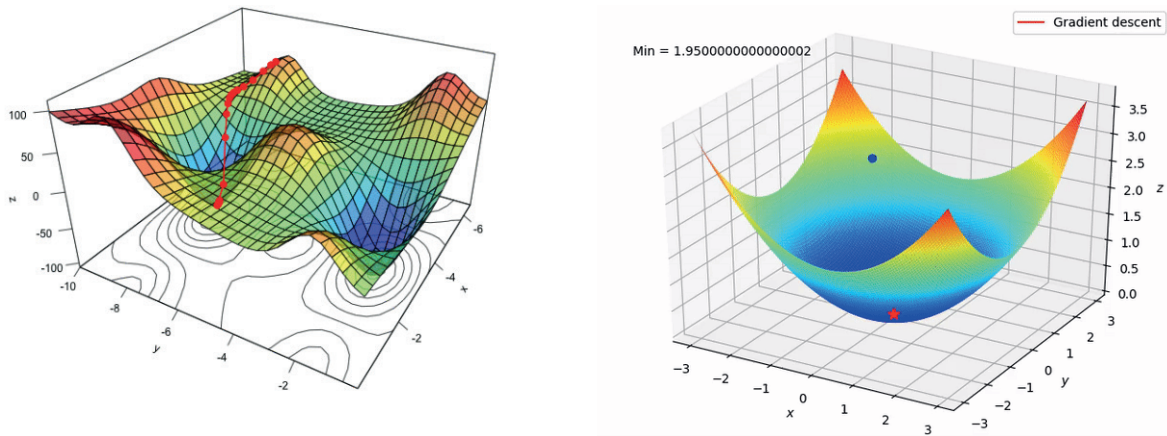
The optimisation problem may not be simple. The graphic below shows two different optimisation surfaces. On the left, we see a well behaved surface where error is minimised easily. On the right, the problem is much more complex, with many local minima, humps and hollows. Both types of optimisation would solve the first problem easily, but an AI/ML based optimisation would be better at the second one.

76. Examples include <https://www.dexfinance.com/>, <https://www.starfetch.ai/>, <https://optionsalgo.ai/>, and <https://quantgatesystems.com/>

77. <https://finscience.com/en/blog/ai-for-fintech-en/ai-for-portfolio-optimization/>, <https://www.linkedin.com/pulse/artificial-intelligence-investment-how-machine-learning-zanev>
<https://www.cfainstitute.org/-/media/documents/book/rf-lit-review/2020/rflr-artificial-intelligence-in-asset-management.ashx>

78. An equation expressing the trade-off of all parts of the portfolio construction process

Image 19: Simple and more complex optimisation problems⁷⁹



Some examples:

The problem of forecast errors in returns or alphas is not well handled by conventional optimisation problems. However, neural networks can be applied which learn which return forecasts create problems. What are known in machine learning as meta-algorithms can be applied. Known as boosting or bagging, these aim to iteratively or sequentially apply techniques which train the model to handle outliers or errors better. Techniques like Adaboost are common.⁸⁰

The objective to be maximum return and minimum risk can be specified in such a way that much more general machine learning approaches can be applied. These are designed to be much more robust to discontinuities like round lots and fixed t-costs and better at handling forecast errors. Many techniques exist, among them:

- Evolutionary algorithms, where the relative “fitness” of each solution is evaluated, leading to a best solution⁸¹
- Particle swarm techniques, where a population of possible solutions (or “particles”) are iteratively shifted through the search space to find a best solution.⁸²

79. Ashwin Singh, <https://medium.com/@ashwin8april/optimization-algorithms-in-deep-learning-4f2c3b53f9f>, Eeswarex, <https://medium.com/@eeswarex/neural-networks-e8393da5c8e>

80. <https://en.wikipedia.org/wiki/AdaBoost>

81. https://en.wikipedia.org/wiki/Evolutionary_algorithm

82. https://en.wikipedia.org/wiki/Particle_swarm_optimization

Summary

As noted above, the difference between an AI model or system and the broader concept of machine learning is blurred somewhat here. We could define it as:

"...a machine's ability to perform the cognitive functions we usually associate with human minds."⁸³

An AI model or system is (just) a type of machine learning (ML). It emulates human behaviour in such a way that it appears intelligent (but is actually only a very sophisticated learning and replication tool). As we currently see them, AIs are much better at automation and efficiency of white collar intellectual tasks than they are at manual labour, but are not cognitive "thinking" machines with imagination and creativity.

Here we discussed the ethics of emerging AI (especially generative AI like chatGPT) and how it might impact human work activities. Ethical considerations abound and the ability of AIs to replace humans in many automated roles is string. We also outlined what an AI model/system actually is, and what it is not.

The final section addresses the main point of our paper – how will investment management be affected? There are many possible impacts, but we have discussed eight areas, across personal finance and institutional investment management. Some of these improve efficiency and financial knowledge, others potentially add value through clever implementation of ideas or through improved trading and portfolio construction.

Ogden Nash, the US humourist poet from the 1950s and 1960s, once said that "Progress might have been all right once, but it went on too long."⁸⁴ We would contend, like Nash, that progress is hard to stop even if we wanted it to.

Appendix; What does chatGPT have to say?

On Friday Sept 1 2023 we asked chatGPT the question: "Describe the impact of AI on investment management in three paragraphs"

The answer was probably as good as any answer we could provide, which suggests that the role of writers of reports like these may be under threat as well.....:

"AI has had a profound impact on investment management, revolutionizing how financial professionals make decisions and optimize portfolios. One of the key ways AI has transformed the industry is through enhanced data analysis and pattern recognition. AI algorithms can process massive amounts of financial data and identify complex patterns that human analysts might overlook. This enables investment managers to make more informed and data-driven decisions, leading to better risk assessment and higher returns. Machine learning models can predict market trends, identify correlations between different assets, and adjust portfolios in real-time based on evolving market conditions.

Furthermore, AI-driven quantitative models have reshaped portfolio construction and risk management. These models utilize historical data to develop strategies that aim to maximize returns while minimizing risk. Portfolio optimization algorithms consider various factors, such as asset correlations, volatility, and historical performance, to construct well-balanced portfolios tailored to an investor's risk tolerance and objectives. This has democratized access to advanced investment strategies, allowing even individual investors to benefit from sophisticated AI-driven approaches that were once only available to institutional players.

However, the integration of AI in investment management also raises challenges and concerns. The potential for algorithmic biases and unforeseen model behaviour poses risks, as seen in instances of market crashes triggered by automated trading algorithms. Additionally, the human touch in investing, such as qualitative judgment and understanding of macroeconomic factors, remains essential to contextualize AI-generated insights. Striking the right balance between human expertise and AI-driven analytics is crucial to harnessing the full potential of AI in investment management while mitigating potential downsides."

83. Ashwin Singh, <https://medium.com/@ashwin8april/optimization-algorithms-in-deep-learning-4f2c3b53f9f>, Eeswarex, <https://medium.com/@eeswarex/neural-networks-e8393da5c8e>

84. From Ogden Nash (The New Yorker, Apr 1959) "Come, Come, Kerouac! My Generation is Better Than Yours". The full quote is "Progress may have been all right once, but it went on too long; I think progress began to retrogress when Wilbur and Orville started tinkering around in Dayton and at Kitty Hawk, because I believe that two Wrights made a wrong."

Disclaimer

Important Information

This material is for general information purposes only. It does not constitute investment or financial advice and does not take into account any specific investment objectives, financial situation or needs. This is not an offer to provide asset management services, is not a recommendation or an offer or solicitation to buy, hold or sell any security or to execute any agreement for portfolio management or investment advisory services and this material has not been prepared in connection with any such offer. Before making any investment decision you should consider, with the assistance of a financial advisor, your individual investment needs, objectives and financial situation.

We have taken reasonable care to ensure that this material is accurate, current, and complete and fit for its intended purpose and audience as at the date of publication. No assurance is given or liability accepted regarding the accuracy, validity or completeness of this material and we do not undertake to update it in future if circumstances change.

To the extent this material contains any expression of opinion or forward-looking statements, such opinions and statements are based on assumptions, matters and sources believed to be true and reliable at the time of publication only. This material reflects the views of the individual writers only. Those views may change, may not prove to be valid and may not reflect the views of everyone at First Sentier Investors.

About First Sentier Investors

References to 'we', 'us' or 'our' are references to First Sentier Investors, a global asset management business which is ultimately owned by Mitsubishi UFJ Financial Group. Certain of our investment teams operate under the trading names FSSA Investment Managers, Stewart Investors, RQI Investors and Igneo Infrastructure Partners, all of which are part of the First Sentier Investors group.

We communicate and conduct business through different legal entities in different locations. This material is communicated in:

- **Australia and New Zealand** by First Sentier Investors (Australia) IM Ltd, authorised and regulated in Australia by the Australian Securities and Investments Commission (AFSL 289017; ABN 89 114 194311)
- **European Economic Area** by First Sentier Investors (Ireland) Limited, authorised and regulated in Ireland by the Central Bank of Ireland (CBI reg no. C182306; reg office 70 Sir John Rogerson's Quay, Dublin 2, Ireland; reg company no. 629188)
- **Hong Kong** by First Sentier Investors (Hong Kong) Limited and has not been reviewed by the Securities & Futures Commission in Hong Kong. First Sentier Investors, FSSA Investment Managers, Stewart Investors, RQI Investors and Igneo Infrastructure Partners are the business names of First Sentier Investors (Hong Kong) Limited.
- **Singapore** by First Sentier Investors (Singapore) (reg company no. 196900420D) and this advertisement or material has not been reviewed by the Monetary Authority of Singapore. First Sentier Investors (registration number 53236800B), FSSA Investment Managers (registration number 53314080C), Stewart Investors (registration number 53310114W), RQI Investors (registration number 53472532E) and Igneo Infrastructure Partners (registration number 53447928J) are the business divisions of First Sentier Investors (Singapore).
- **Japan** by First Sentier Investors (Japan) Limited, authorised and regulated by the Financial Service Agency (Director of Kanto Local Finance Bureau (Registered Financial Institutions) No.2611)
- **United Kingdom** by First Sentier Investors (UK) Funds Limited, authorised and regulated by the Financial Conduct Authority (reg. no. 2294743; reg office Finsbury Circus House, 15 Finsbury Circus, London EC2M 7EB)
- **United States** by First Sentier Investors (US) LLC, authorised and regulated by the Securities Exchange Commission (RIA 801-93167)
- **other jurisdictions**, where this document may lawfully be issued, by First Sentier Investors International IM Limited, authorised and regulated in the UK by the Financial Conduct Authority (FCA ref no. 122512; Registered office: 23 St. Andrew Square, Edinburgh, EH2 1BB; Company no. SC079063).

To the extent permitted by law, MUFG and its subsidiaries are not liable for any loss or damage as a result of reliance on any statement or information contained in this document. Neither MUFG nor any of its subsidiaries guarantee the performance of any investment products referred to in this document or the repayment of capital. Any investments referred to are not deposits or other liabilities of MUFG or its subsidiaries, and are subject to investment risk, including loss of income and capital invested.