

Cambridge Judge Business School

Artificial Intelligence & Machine Learning in Finance

IS IT JUST ANOTHER BUBBLE?

Conference summary

Centre for
**Endowment Asset
Management**



**UNIVERSITY OF
CAMBRIDGE**
Judge Business School

Supporters

This event was organised by the Centre for Endowment Asset Management (CEAM) at Cambridge Judge Business School (CJBS). We would like to thank FTSE Russell, Invesco Ltd. and Newton Investment Management, without whom this event would not have been possible. The philanthropic support these organisations provide to CEAM enables us to facilitate research engagement and dissemination activities, like this event, amongst academics and practitioners.

With additional thanks to Invesco Ltd. for its support and collaboration in producing this event summary.



Contents

Foreword 4

Introduction 5

Insights and Interpretability 6

Knowledge extraction: from theory to practice..... 7

Humans or machines – or both? 8

Consent, commitment and competition..... 9

Foreword

The rise of artificial intelligence (AI) and machine learning (ML) techniques represent one of the most significant current trends in finance.

On 8 October 2019, Cambridge Judge Business School's Centre for Endowment Asset Management hosted a conference exploring how these techniques are being applied in the financial sector at present and where they might lead in the future.

The conference, entitled "Artificial Intelligence and Machine Learning in Finance: Is It Just Another Bubble?", attracted over 70 asset managers, asset owners, academics, entrepreneurs, regulators and policymakers. With the focus on the impact of machine learning (ML) and artificial intelligence (AI) applications in the financial sector, we discussed topics such as "black box" concerns, the likely scope and limitations of AI and ML, the delicate balance between human and machine inputs and the ever-growing relevance of regulatory and ethical considerations – most notably in relation to data privacy.

On behalf of the Centre for Endowment Asset Management, I'd like to thank FTSE Russell, Invesco Ltd. and Newton Investment Management for their valued support for this event and to thank the speakers and guests who gave up their time to join us and contribute to the event.

Sarah Carter
Executive Director
Centre for Endowment Asset Management
Cambridge Judge Business School

Introduction



Many of the ideas that underpin artificial intelligence (AI) and machine learning (ML) first emerged decades ago. Alan Turing presaged a computer's ability to cope with a "combinatorial explosion" when he outlined the basics of a chess-playing program in the late 1940s. Herbert Simon and Alan Newell's original "thinking machine", which used heuristics to replicate rational human thought, was unveiled in the mid-1950s. The Digital Equipment Corporation's XCON system, which saved the company millions of dollars by automatically selecting computer components based on a customer's requirements, was first deployed in 1980.

So, what is different today? We now have unprecedented processing power. There is a superabundance of data. Perhaps above all, the domain of AI and ML is no longer the exclusive preserve of a few trailblazers: it has become a familiar element of almost everything that we do in our professional and personal lives.

In the sphere of finance, an industry that tends to absorb novel technologies especially quickly, the impact is proving transformative – even disruptive. As is invariably the case when a sector embraces significant innovation, opportunities and challenges are appearing in roughly equal measure.

At our recent conference, *Artificial Intelligence and Machine Learning in Finance: Is It Just Another Bubble?*, we attempted to take stock of where things stand today and where they might go from here. Drawing on the insights of a diverse array of stakeholders from around the globe, we explored issues such as "black box" concerns, the likely scope and limitations of AI and ML, the delicate balance between human and machine inputs and the ever-growing relevance of regulatory and ethical considerations – most notably in relation to data privacy. Although the event was conducted under the Chatham House Rule, which places obvious limitations on any report of proceedings, we felt that it was important to try to encapsulate a lively and informative series of presentations and discussions – hence this summary.

Whether AI/ML in finance will turn out to be another bubble remains a moot point. Most of our conference participants envisaged its continued rise and influence; others doubted its capacity to alter the landscape beyond recognition; all agreed that it must be kept under control. In the absence of a definitive answer, maybe there are three points that are particularly worth bearing in mind.

First, what we have at present, even after all these years, is not artificial intelligence: very strictly speaking, it is *augmented* intelligence. Second, AI and ML still rely on humans to call the shots. Third, finance is a world in which the rules are always changing.

The story of AI and ML, whether in this or any other setting, is clearly far from over. We hope that this report sheds at least some light on the realities, the possibilities and the potential implications of this still-unfolding phenomenon.

Dr Pedro Saffi
Reader in Finance, Cambridge Judge Business School and Academic Fellow, Centre for Endowment Asset Management

Insights and Interpretability

"We're trying to develop an idea of what we're getting out of these black boxes."

Key points at a glance

AI and ML can provide investment insights that more traditional approaches might not reveal.

Research is increasingly directed towards understanding exactly how these insights are generated.

Simplicity and interpretability are crucial to addressing AI and ML's "black box" challenges.

The threat of "style drift" might need to be considered when moving towards AI/ML-led models.

Investors like to understand how returns are generated. Even the most agreeably profitable strategy demands an explanation, otherwise its success might just as well be ascribed to chance. With AI and ML playing an ever-greater role in the search for alpha, unravelling the "black box" mysteries that new methods can entail presents a significant challenge.

A growing body of literature highlights the effectiveness of using AI and ML to predict and explain stock returns. Flexible, non-linear approaches can be especially helpful in illuminating what one researcher described as "dusty corners" – those areas "where things look different and it's important to capture the uncertainty".

The conference featured a demonstration of an ML-driven, non-parametric regression model that incorporates economic information and time variability. Analysing more than 600 months' worth of data, this model consistently identified short-term reversals as key to explaining the non-linear relationship between predictors and equity returns during the period from 1963 to 2015.

The appeal of models that are able to provide such insights is obvious, but what should their essential attributes be? As one practitioner observed, the most common Achilles' heel of such innovations is their sheer complexity: the more intricate they become, the less likely they are to address the vital "black box" question. It is imperative that they bring clarity, not obscurity.

One specialist proposed a number of stipulations for an AI/ML-led model: it should be as simple as possible; it should be capable of coping with "noise"; it should be powerful; and, crucially, it should be interpretable. "You have to be able to explain its performance to the client base," he said. He added that most current strategies would not satisfy all four provisos.

Another practitioner sounded a similar note of caution, remarking that AI and ML approaches bring "promise and pitfalls". He told the conference that his own real-world portfolio simulations had clearly underlined the value of limiting complication. "The more constraints you have," he said, "the harder it is to exploit the signals that machine learning generates."

An additional difficulty could lie in the threat of "style drift". Given that conventional quant investing represents a crowded space, more investment firms might be tempted to "morph" their products and solutions. With a rise in AI/ML-related terminology in client-facing literature suggesting that such a shift may already be underway, one practitioner warned: "AI is very different from factor-driven, traditional quant strategies. There could be a risk in terms of asset allocation as a drift occurs."

Knowledge extraction: from theory to practice

“As HL Mencken said: ‘For every complex problem there is an answer that is clear, simple and wrong.’”

Key points at a glance

Combining big data and formative rules is key to the quest for ongoing “knowledge extraction”.

This approach is already being used to predict human behaviour in a range of settings, including finance.

There are important challenges to consider, among them ethical implications and “black box” concerns.

AI and ML’s influence nonetheless seems likely to spread – possibly to areas such as regulation.

What might be called “classical” AI tends to operate on a relatively straightforward, heuristic-led, “If... then...” basis. This can produce a kind of first-order logic whereby something that is demonstrated might be deemed proven in perpetuity. By contrast, advanced ML uses neural networks – algorithms loosely modelled on the human brain – to exploit a vast and ever-growing amount of data to develop powers of reasoning that are much greater and whose outcomes are more easily explained.

This process thrives on one of the defining phenomena of our age. “Big data is a valuable fact of life,” one researcher told the conference, “and that’s why this isn’t another a bubble.” Yet it is not sufficient for ML merely to accrue a near-boundless wealth of information: it has to make sense of it as well, and for this to happen there have to be rules.

By way of illustration, consider the Waymo autonomous vehicle. It can draw on approximately 10 billion miles’ worth of driving data, yet it would most likely struggle to survive a short cruise around central London without a firm grasp of the Highway Code or a similar set of guiding principles. Continued “knowledge extraction” relies on a machine being able to constantly revise its understanding – not only by following rules but also, over time, by recognising where exceptions might be made.

The conference heard how this approach is now being used to predict human behaviour in a variety of settings. In the sphere of online gambling, for example, ML is able to detect patterns of play that indicate when a customer might be at risk and should take a break. Such applications are important from numerous perspectives – including regulatory requirements, system performance and improving

consumer confidence – and can allow different stakeholders to make different interventions. So what about the opportunities and challenges around knowledge extraction in the financial services industry?

The deployment of AI and ML has become a day-to-day issue for many of the largest entities in the sector, with ethical implications and “black box” interpretability among the main concerns. Outlining two hypothetical situations, a representative of a global investment bank asked: “Should we base a decision about a mortgage application on data that machine learning might highlight but which the applicant may not volunteer personally? And is it within our fiduciary duty to invest in artificial intelligence if we see a 12% growth in profit from it but no-one knows why?”

One practitioner revealed that his company, which uses a “pure AI play” approach to investment, mostly employs tech experts with no experience of finance. “This is intentional,” he said. “We don’t want them to be rooted in 50 years of economic thinking.” He also suggested that the advent of quantum computing, as well as curtailing the energy-intensive processes often used today, would most likely further extend AI and ML’s influence.

Could AI and ML one day help financial institutions deal with regulatory change? And, if so, would machines be capable of making sense of a raft of requirements while at the same meeting the desired ethical standards? Several speakers remarked that progress in this respect has improved substantially during the past few years. “The real trick is to link all this to control systems,” one researcher told the conference. “That’s a connection that needs to be made.”

Humans or machines – or both?

"I think this presents a real opportunity for us to put our best selves into decision-making."

Key points at a glance

Striking an optimal balance between human and machine inputs will be key to the future of financial services.

Employees might fear the ever-greater adoption of AI and ML, which can also be hampered by legacy issues.

Despite such hurdles, these innovations present unprecedented opportunities to make "human" decisions.

Given the nature of the markets, the prospect of an entirely machine-led financial services industry is unlikely.

Conducted more than 70 years ago by Alan Turing, what is now widely thought of as the first-ever study of AI set out to determine whether machines can do what humans can do. The focus has long since shifted to what machines can do that humans *cannot* do, yet it is just as important to recognise the opposite – that is, what humans can do that machines cannot do.

The need to find an ideal balance in this regard was the subject of one of the conference's liveliest discussions. Some contributors posited that the AI/ML revolution has barely begun, while others argued that it might turn out to be far less transformative than generally expected. Amid a diversity of views, however, one point of consensus emerged: humans and machines alike will continue to have valuable parts to play in the world of finance.

The conference heard how the introduction of AI and ML has prompted resistance in a variety of settings. At a major news-gathering organisation, for instance, journalists were reluctant to concede that a program that enhanced the detection, verification and publication of breaking news from Twitter – thereby giving the company a substantial competitive advantage – had augmented their own work. Drawing on personal experience in a number of industries, one specialist identified employees' reservations and an attachment to existing technology as the biggest human-related problems associated with the adoption of AI and ML.

Yet the optimisation of this technology still necessitates human inputs of many kinds. As one practitioner at the cutting edge of developments in this field observed: "As a human, I want to do the right thing – and if we want AI to have the same

attitude then we need to put it in. Remember that machines don't know what we're hoping for. Everyone in the life-cycle of AI – designers, coders, overseers, users – needs to ensure that we're making 'human' decisions."

According to a leading tech entrepreneur, a superior data infrastructure is another imperative if human-machine interaction is to function to best effect. Making the case for a more "atomic" approach, he warned of the threat of a "data oligopoly". "AI shouldn't be a big thing that sits on a big server," he told the conference. "It should be a small thing that sits in each asset. We still don't really know how to collect the data and structure it. We need to make it atomic rather than having something that's huge and autonomous."

The most forthright rejection of fears that AI and ML could change finance beyond recognition – and at humans' expense – came from the CIO of a global investment manager. He acknowledged that the use of AI is growing, even among active managers, but insisted: "I'm not convinced that it's going to take over the world just yet."

What should be remembered, he suggested, is that the financial markets are "a social construct that doesn't have immutable laws". "Most of the things that have blown up resulted from the extrapolation of trends that ended abruptly," he said. "There are many dynamics where things turn on their heads. There's huge scope for AI, but no amount of pattern recognition will work if the patterns don't repeat."

Consent, commitment and competition

“Algorithms require a huge amount of data, but does all that data come with consent? This is something we should all be concerned about.”

Key points at a glance

With the gathering and processing of data now considered a norm, consumers are falling prey to “consent fatigue”.

Many value their privacy yet have little concept of how much information can be gleaned from an online presence.

Consumers, providers and regulators alike must ensure that data use reflects a broader commitment to the greater good.

As the wider understanding of this issue improves, customers are likely to favour organisations that prize transparency.

The superabundance of data that acts as the lifeblood of AI and ML undoubtedly brings benefits. Most obviously, organisations that have a wealth of information about the individuals they serve should be able to tailor and personalise products and services to an unprecedented degree. But the data explosion also brings challenges, foremost among them those around consent and privacy.

One problem is that the processing of data has become so fundamental to consumer activity that many of us have succumbed to “consent fatigue”. “Most people claim their data privacy is worth a lot, but then they act as if it’s worthless,” a leading researcher in this field told the conference. “They’re reassured by the presence of legal text when they go online, even though they have no idea what it actually means.”

According to a study of the data privacy contracts used by almost 3,500 US firms, even the simplest text of this sort requires years of formal education to comprehend in full. The policies of large firms tend to be longer yet easier to find, while those of smaller and less technologically sophisticated firms tend to be notably lacking in clarity. Whether a company process data in-house or shares it with a specialist intermediary depends on factors such as opportunity cost, value and litigation risk.

In addition, many consumers do not appreciate the depth and complexity of the inferences that can be drawn from an online presence. Another researcher explained how big data has been used to predict personality traits, political views, intelligence levels and other attributes more accurately than a person’s partner, friends or colleagues could. Our own activities and preferences can even be gauged

based on those of the people with whom we connect electronically.

Considering this reality, data privacy should be seen as a core element of a commitment to act in the interests of society as a whole. “People have become more aware of their fundamental rights over privacy,” a lawyer for an international organisation said. “Firms need to be reminded that customer-facing tech has to be compliant and in the best interests of the consumer.”

The regulatory community is also taking a keen interest in how big data is reshaping the interactions between consumers and providers. As one representative explained: “We have enormous predictive potential, incorporating information that we didn’t have even five years ago. In many ways this is great if you’re in a good relationship – but not if you’re in an exploitative relationship.” The conference was told that regulatory bodies are well aware that various sectors – not least financial services – are being confronted by “new questions”, including concerns around “‘black box’ approaches and absence of explainability”.

Studies have indicated that larger companies are usually less likely to use data in ways that might make their customers uncomfortable. Smaller enterprises are more willing to use data as they please, as are a few giant corporations that are not overly vulnerable to competition. On balance, the conference heard, data is currently employed in more positive ways than negative ways – but what of the future? “Consumers should ultimately choose companies that are transparent,” said one researcher. “Transparency will be a competitive advantage.”

Centre for Endowment Asset Management

Cambridge Judge Business School
University of Cambridge
Trumpington Street
Cambridge CB2 1AG
United Kingdom

ceam@jbs.cam.ac.uk

www.jbs.cam.ac.uk/ceam

