

Ashridge Strategic Management Centre

Members Meeting

6th March 2019

Minutes of Meeting

In attendance

Paul Barrett	Babcock International
Philip Meyers	ABF
Lydia Ogilvie	National Grid
Adam Prince	Shell
Stein Rasmussen	SBM Offshore
Sarah Sandle	Rolls Royce
Ben Slater	BP
Paul Titcomb	Atkins
Tara Yamashita	BAE Systems

From Ashridge Strategic Management Centre

Felix Barber
Stephen Bungay

Guest Speakers

David Dean, Axiata, WEF Global Future Council on Digital Economy and Society
Angie Ma, Faculty.ai

Felix Barber, Angie Ma and David Dean: AI and Strategy

Felix introduced the subject by quoting Google CEO Sundar Pichai's statement that artificial intelligence is 'probably the most important thing that humanity has ever worked on'. The meeting focused on what it means for member companies.

AI has been defined as 'the science of making computers do things that require intelligence when done by humans' and as 'a computer solving a problem that it has not explicitly been told how to solve'. Felix suggested that it can be understood as a combination of three things:

At its core is **machine learning**, whereby a model is given lots of examples for it to learn from. It is programmed with a learning algorithm

which allows the machine to learn for itself, unspecified rules to allow it to define how to solve a problem and large data sets to allow it to learn from a large number of examples.

The second thing is **sensing**, which enables the computer to gather data for itself.

The third thing is **taking action**, whereby the computer draws its own conclusions from the data it processes and take autonomous action.

These three things are brought together in the context of human governance.

Angie Ma: Improving Business with AI

Angie focused on the central theme of machine learning, drawing on her experience of over 300 projects Faculty.ai has worked on with clients, covering 23 sectors in 8 countries.

Whilst it has been talked about and developed for decades, the actual deployment of AI has really taken off over the past two years. It is predicted to become as basic as steel and enable things we cannot predict. But currently we are building a hammer not a space station.

Angie believes that AI will develop in particular ways.

Commodity AI, based on huge data-sets, will be owned by the technology giants like Google, Amazon and Microsoft. They feed off horizontal businesses which provide them with data, and are unlikely to go into depth around any set of vertical variables.

However, the real impact of AI will come from its wide adoption by other companies, all of which specialise in something. They will use deep data rather than big data, and will need to make choices about what to buy in as a 'black box' and what to build themselves in order to retain competitive advantage.

To do this, companies need to hire data scientists, decision scientists and AI architects. All of these are rare skills, and such people are also difficult to retain, so their availability may act as a constraint on growth. Faculty.ai began as a fellowship to train and transition PhD and post-doc scientists into professional data scientists. Computation will remain fundamental to data science and coding will remain fundamental to computation.

There were a few comments about where AI could have value and the barriers to applying it:

- Using enzymes in pulp and paper production can reduce costs by 5%. Manufacturers are conservative about using them because of the risk of paper breakage. Using cheap sensors would allow them to isolate what leads to breakages, but most players have given up trying because it is so hard;

- The maintenance of gas turbine engines could be optimised by applying AI to the large amount of data available, but the problem is that the data needed is held by different companies, such as manufacturers and airlines, and it is not even always clear who owns it.

However, there are examples of AI being successfully deployed to optimise processes. Russian steel manufacturers are using it to identify the precise make-up of scrap metal and so optimise the smelting temperature. It is also potentially of great value for running low-cost experiments in FMCG businesses, all of which want to understand human behaviour.

In general, the conditions under which AI can add real value to a business are when there is some combination of the following factors:

- a high data volume;
- the data is highly dynamic (i.e. fast changing);
- not all the relevant variables are identifiable in advance;
- there is a high value to creating more precision;
- there is value to exploring a large number of scenarios.

The actual value extracted depends on the design of the model.

How much data is needed depends on the application. Some need specific domain knowledge. There is some competitive risk that companies can take data from others and monetise it. Google maps, for example, takes data from car manufacturers.

Early examples of computers out-performing humans were restricted to very narrow domains with set rules, such as playing chess. IBM's Big Blue beat Garry Kasparov in 1997. In 2017 Google's Alpha Zero taught itself chess in 24 hours and consistently beat the best conventional computer chess programme, Stockfish 8, in all of 100 games. However, in complex cases without set rules in which the variables are unknown, only AI will work, and the results of analysis using AI can be surprising. For example, a Chinese company is doing credit rating based on how you unlock your smart phone.

Angie went on to take us through some case examples.

Her company built AI software using image and audio recognition to help the Home Office to identify terrorist videos on the web. The software was trained using images chosen by humans. The front-end work was extensive, with only obvious cases being used in order to avoid human bias. Once running, the system needs constant checking, but this only requires 20 people instead of hundreds. Google had claimed that the task was too hard for them, but it was

probably more likely that they had no interest in taking it on because they are a platform which does not want to accept responsibility for content.

Another case was detecting fraud for a large internet payments company. Though only 0.0001% of transactions were fraudulent, the cost of fraud and of detecting it were high. If it experienced too much fraud, the company could have its licence revoked, but false positives were also a serious problem. The existing system required 200 rules and was still inaccurate. The model identified millions of variables, detected fraud with 99.9% accuracy and reduced false positives. The fraud detection team has been reduced from 200 to 20, saving about 20% of the company's overall salary budget.

The main obstacle to finding enough relevant data is the way it is organised. Even if a company owns all the data needed, it is often in different datasets. A start-up with no legacy has an advantage here.

Members shared some of their experiences in using AI:

- In building water treatment plants we now give the output required to a programme and it is designed using AI;
- We are starting to use AI in matching electricity supply with demand. Hitherto, it has been done by very skilled people, but when they leave, they take their knowledge with them;
- We have enough data to do predictive maintenance of military equipment. It is hard to apply AI to the rest of the business, where OEM's have the data;
- The maintenance of floating production systems is costly and dangerous, so more precise prediction would be valuable. We use a 3rd party. One concern is losing the value of the data we have.

Felix Barber: AI Based Strategic Transformation

Returning to the model he introduced at the beginning of the meeting, Felix explained that innovation is taking place in sensing, with for example, the use of drones, and acting, with the use of robots. Some things which are simple for humans, such as lacing shoes, are difficult for robots, but it is easier to share data between machines than between humans. A lot of disruptive AI exploits this.

For example, a self-driving car uses many mechanisms such as cameras, radar, lidar, ultrasound etc. to generate data, takes data from GPS and 3D maps to learn, and uses autonomous steering braking and accelerating to act.

Breaking things down into these three categories both helps to understand where the value lies and also to identify vulnerabilities.

Felix introduced a model derived from the Lean Canvas (which is itself a modification of the Business Model Canvas) to help identify opportunities. He illustrated it with the example of Blue River Technology, founded in California in 2011, which uses AI to help lettuce farmers.

The problems faced by lettuce farmers include thinning lettuces and deciding which ones to grow, and then controlling weeds using herbicide. The former is labour-intensive and the latter expensive, bad for the environment and increasingly ineffective as weeds become resistant to blanket herbicides. Blue River has developed a 'See and Spray' technology which uses sensors to distinguish lettuces from weeds, and spray herbicide directly onto the weeds, saving 90% of the volume used in traditional blanket spraying and allowing different sprays to be used. They originally developed a model that could sense and learn, but had no means to act. They sold the business to John Deere in 2017. Deere built the spraying equipment, thus completing the full AI value chain. If this technology spreads it will cause a problem for herbicide manufacturers, who will see their volumes drop dramatically.

Felix asked the group how they evaluated AI opportunities.

One member reported that they often lacked a scarce resource held by an incumbent, and were recipients of elements they needed to outperform. They have a portfolio of ventures.

Another suggested partnering with companies like Faculty.ai, who sometimes provide off-the-shelf products and sometimes tailor-made solutions. One problem in seeking partners is that it is hard to know just what specialism you want. Members were reminded that Felix' book 'Collaboration Strategy' addresses that problem.

David Dean: Responsible AI

AI is in the public eye and raises controversial issues around reliability, privacy, transparency, inclusiveness, trust, values and ethics. 72% of US business leaders believe AI will define their future and up to 76% are concerned about stakeholder trust and the potential for bias. But consumers are even more skeptical – 96% would not trust AI in hiring employees and many are concerned about its effect on jobs, trust and indeed the future of humanity.

David suggested that dealing responsibly with AI involves three things: business goals and ethical and legal constraints; governance; and engagement.

In setting goals, businesses have to make trade-offs. Should a call centre optimise productivity or customer loyalty? Do callers have the right to know if the voice is non-human? There is no legislation about this.

Algorithms learn from data and therefore reflect any biases within it. Microsoft's 'Tay' chatbot launched in 2016 turned into a misanthropic Nazi in 24 hours, and a Google search of 'beautiful hairstyles' exclusively features young white women.

Ethical issues encapsulated in the 'trolley problem' are particularly difficult because there are no universally accepted views on 'the right thing to do'. MIT have set up a website called Moral Machine to try to illuminate ethical issues by using crowd sourcing.

Manufacturers of driverless cars must programme them to spare pedestrians or passengers in an emergency. In Japan, there is a strong emphasis on sparing pedestrians, in China an equally strong one in favour of sparing passengers. In countries with a more individualistic culture there is a bias towards sparing the young, in more collectivist ones towards sparing the elderly. This creates a problem in building a global model. Even when survey participants claim to prefer a car that spares pedestrians, they would prefer not to ride in them.

One observation from the floor was that companies face lots of ethical issues, many of which are more important than those posed by AI. David replied that AI raises new issues of data control which companies are not set up to deal with. For Facebook, the issue is existential.

David suggested that companies needed to extend the scope of their governance to create a clear accountability framework, and audit and respond to public responses. AI issues also needed to be included in risk management processes and a few, like Salesforce, have started employing a Chief Ethics Officer.

The other aspect of deploying AI responsibly is engagement with employees and the public, both through public debates and regulators. Nokia is training all staff in AI, starting with the Chairman, and the Finnish government is planning to give all its citizen lessons in AI, starting this year with the first 1%. Microsoft has developed a set of principles specifically for AI: fairness, reliability and safety, privacy and security, inclusiveness, transparency and accountability.

We ended the afternoon with summary comments and questions from Members.

Round the table comments

- Working with an algorithm is not a sell and forget process, it has to be maintained. Using third-party suppliers is problematic. What if they get acquired?
- Is AI more than another form of automation? I am sceptical that it is qualitatively different.
- The distinction between commodity and specialised AI is useful. The other issue that concerns me is where governance should rest.
- I was particularly interested in the Blue River example because we are involved in arable farming. It is very high-tech and in some ways the canary in the coalmine.

- It bolsters the value of data processing in thinking about the demand for power. We are highly regulated, and AI can in fact help us to deal with things like safety and reliability.
- How to organise for all of this could be a theme for the strategy department.
- We need to learn how to do this. It often comes over as a technology product, but the skills and education side of it are very important.
- We are looking at how to use it in the business. You need middle managers who know enough to judge where the value lies. But I do share some skepticism – is it just incremental change or does it herald a very different future?
- It is developing at very high speed. Many people do not realise just how much our current society is underpinned by data, and this is changing how that data can be used.

Future meetings

The next Members' Meeting will be held from 13.00 – 17.30 on 15th May 2019. It will take place at the same venue, The Royal Horseguards Hotel, 2 Whitehall Court, Whitehall, London SW1A 2EJ.

The topic will be 'Navigating in Uncertainty' and will be led by Stephen Bungay and Rebecca Homkes.

Members are also reminded about our **Seminar** to be held at De Vere Venues on **16th May**, when the subject will be Organisation Design and Collaboration and will be led by Andrew Campbell.

Please note that **the seminar following that one**, which is about strategy in uncertainty, will take place on **July 25th**, not July 11th as previously advertised.

Members are also reminded that our **Strategy Bootcamp** will run again from **September 30th to 3rd October**. If they book places for members of their team or other executives before 14th June, they will benefit from an early bird discount. Some of you have already done so. A brochure containing the full details will be sent out to all Members shortly.