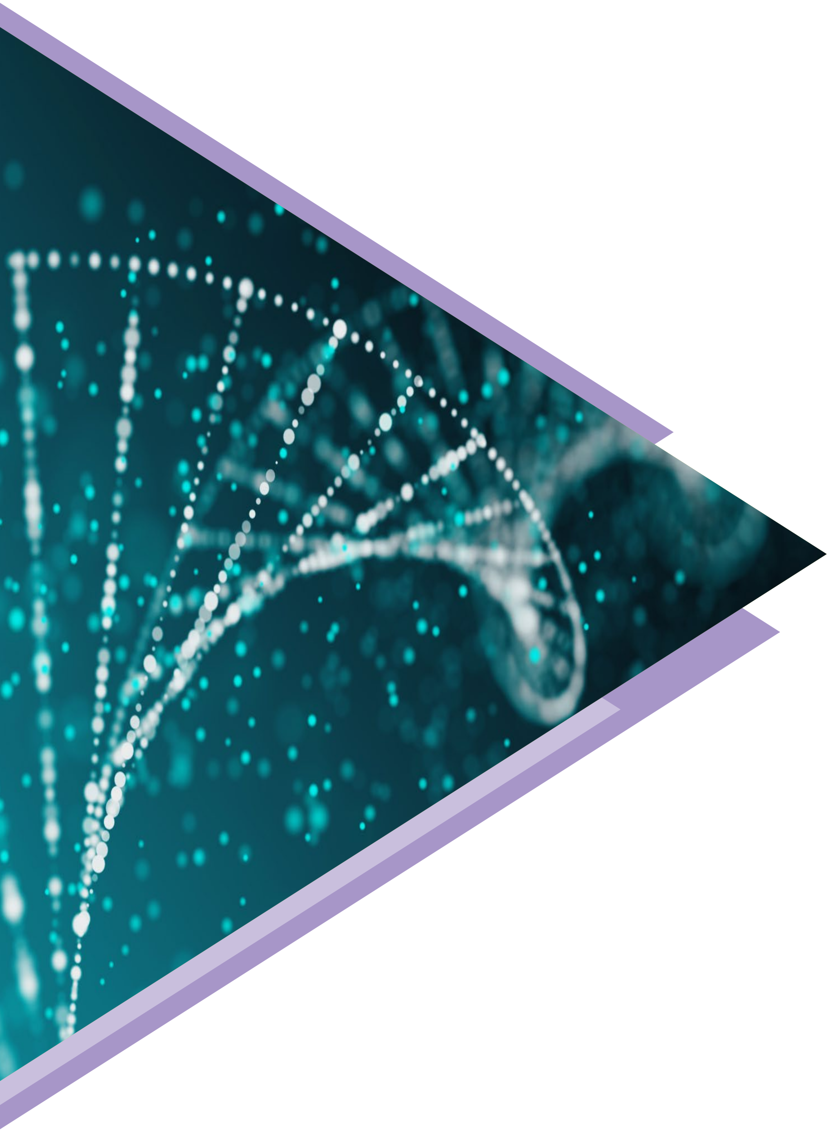

An academic perspective:

AI'S LEAP FORWARD, HUGE POTENTIAL AND UNDOUBTED LIMITATIONS

Advances in computing power and data availability have accelerated AI's evolution and it is now entering our daily lives. Even so, it is still in the foothills of the Himalayan task of developing systems with deep understanding. Interview with Professor David Barber, Director of the UCL Centre for Artificial Intelligence and fellow of the Turing Institute.

QUICK READ

- **AI suffers from a blurring of the boundaries – it's about machines being able to mimic the way humans work, rather than simply the analysis of large data sets.**
- **Computing power and data availability have combined to enable a period of accelerated development in machine learning, a major data-driven sub-field of AI.**
- **AI works best when carrying out limited and well-defined tasks where large amounts of data are available to train the algorithms effectively.**
- **While the prospect of "artificial general intelligence" remains distant, the real-world applications of AI are going to be hugely economically significant.**

Professor David Barber
Director of the UCL Centre for Artificial Intelligence

Our everyday lives are increasingly spent interacting with technology that replicates abilities. Advances in AI have resulted in voice recognition software that allows us to give instructions to Siri or Alexa. They have created sophisticated predictive text functions in email programs, online customer service chatbots and the telephone-based systems now being deployed in call centres, all of which depend on natural language processing. Translation tools and digital assistants that are capable of turning speech into text work in the same way. Image recognition software, as deployed in facial or number plate recognition systems and autonomous vehicles, also represent everyday examples of AI in action.

Innovations in robotics, a closely related field that offers exciting possibilities in areas such as driverless cars, warehouse automation and personal care for the elderly or infirm, also depend heavily on advances in AI. These machines all use AI to replicate humans' ability to interpret and interact with the physical environment, as well as drawing on insights from neuroscience into how humans function.

The emergence of AI as an increasingly common feature of modern life suggests we are on the cusp of a transformation that will produce vast changes in the way humans live and work. However, to interpret and navigate the effects that AI is likely to have on society and the commercial world, we must explore how and why it appears to have undergone such a major leap forward in recent years, and appreciate its current limitations as well as its undoubted potential.

Why are we talking about AI now?

Dr David Barber, Professor of Machine Learning at UCL and Director of the UCL Centre for Artificial Intelligence, points out that attempts to create human-like abilities in man-made systems date back centuries. Indeed, Prof Barber is a fellow of the Turing Institute, which recognises the pioneering role of Alan Turing, who died in 1954, in the development of the discipline. Turing and fellow mathematician and economist David Champernowne wrote their groundbreaking chess program, Turochamp, in 1948 during their research into AI. But the algorithm that powered Turochamp was too complex to run on the computers of the time and Turing was only ever able to execute the program manually, using paper calculations.

Computing power's unstoppable rise

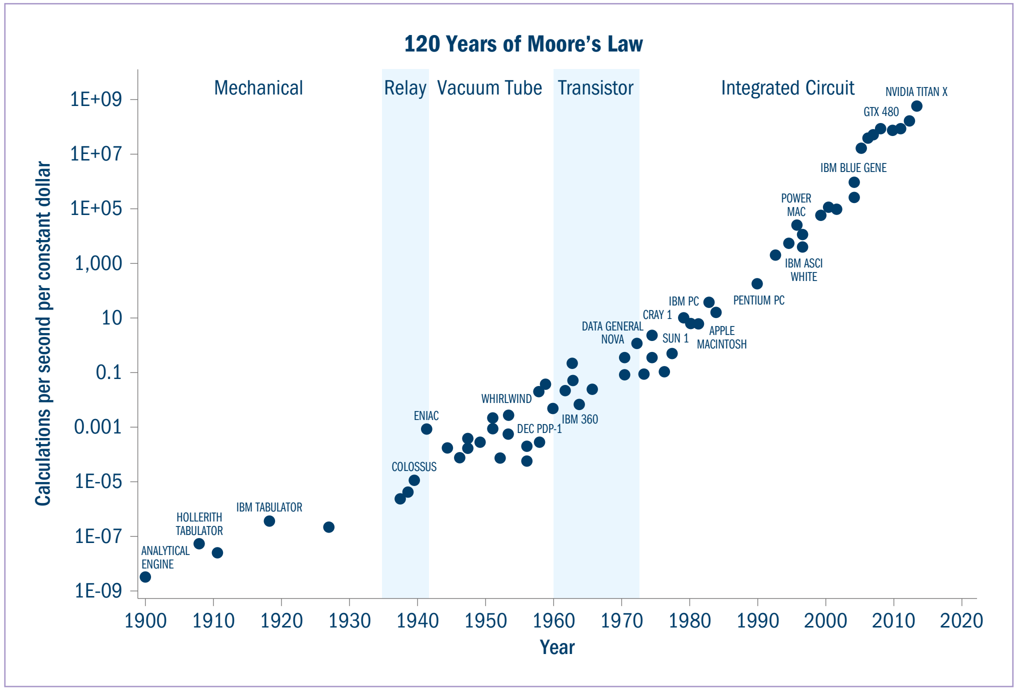

Source: As of 2018. https://www.britannica.com/technology/Moores-law.

Information and opinions provided by third parties have been obtained from sources believed to be reliable, but accuracy and completeness cannot be guaranteed. The information is not intended to be used as the sole basis for investment decisions, nor should it be construed as advice designed to meet the particular needs of an individual investor.

This anecdote illustrates an important point. The basis of many of the algorithms in use today are not new. What has unlocked their potential and, therefore, that of AI, is the abundance of computing power that has become available in recent years as processing speeds have increased. An advanced image recognition system takes about a week to train using today's single NVIDIA GPU-based computers. Performing the same number of calculations using the best workstations available in the early 1990s would have taken hundreds of thousands of years. Gains in computing speed, compounded over decades, have delivered hardware capable of allowing AI to operate in real-time.

The second vital factor in the emergence of AI has been the increasing availability of data. As the quantities of digital data created and stored have multiplied rapidly over recent years, so data sets large enough to train algorithms to high levels of accuracy and proficiency – for example, images for use in teaching object recognition – have been created.

Machine learning emerges as the dominant approach

Taken together, these two factors – computing power and data availability – have combined to enable a period of accelerated development in machine learning (ML), a major, data-driven sub-field of AI. As a result, over the past 15 years or so, ML has become the dominant paradigm within AI and is largely responsible for the advances that underpin the applications we are most familiar with today.

Within ML, a key avenue of development dating back decades involved neural networks – systems loosely based on the structure of the human brain. Having been largely out of favour for many years, neural networks became the key focus of ML research once again in 2006, when a small group of researchers demonstrated that with access to enough computing power, the technique offered significant improvements in the results obtained.[1] Other major advances followed quickly. Soon afterwards, researchers were able to adapt graphics processing units (GPUs) developed for computer gaming to speed up the process of training ML algorithms 100-fold.

Armed with rapidly improving technology tools, a research group led by Geoffrey Hinton made a breakthrough in image recognition using ML in 2012 and was immediately acquired by Google. Hinton's team then went on to quickly create a speech recognition system far ahead of any previous system. ML and its variants, such as deep learning, had become the key technique within AI.

Deep Blue, AlphaGo and the limits of gameplaying

The best-known landmarks in AI tend to be moments such as the victory of IBM's Deep Blue computer over world chess champion Gary Kasparov in 1997, or of Google-owned Deepmind's AlphaGo over Korean Go champions Lee Se-dol in 2016 and Ke Jie in 2017. These highly symbolic events naturally appeal to our fascination with the idea of machines that can outwit humans. But how significant are they, in fact?

Prof Barber observes: "A story that ML researchers like to tell is that we had a machine that could beat the best human chess player in 1997, but we still don't really have a robot that can smoothly and reliably pick up a chess piece and move it." High-profile achievements like these are important to generate attention, but in pure research terms they have turned out to be far less significant than many assume, he argues. "What's ultimately important is not the ability to play chess or Go, but the delivery of systems that will be useful in our daily lives. The rest is largely entertainment."

The most important challenge for AI, he argues, is moving it beyond the enclosed, rules-based world of games and making it good enough to operate alongside humans in the much more complicated environment of our daily lives.

"What's ultimately important is not the ability to play chess or Go, but the delivery of systems that will be useful in our daily lives. The rest is largely entertainment"

[1] One particular milestone is reported here https://science.sciencemag.org/content/313/5786/504 in which neural networks were shown to vastly outperform traditional methods for image compression.

> "Speech recognition is a good example where we have quite good performance now. But it's still very superficial – the machine doesn't actually understand in a deep sense what you're saying"

Where today's AI works best

If the research significance of AI systems that can beat chess or Go champions may be overstated, however, in one respect perhaps these achievements do carry some wider significance. Games such as Go or chess are extremely complex, rules-based problems for which very large quantities of game-play data are available that the ML algorithms can learn from. A newer version of AlphaGo, AlphaZero, taught itself to play Go, chess and shogi to a higher level than its predecessors simply by applying the rules it had been given to play practice games, thus removing the need for data from completed human games.

It is no coincidence that the real-world fields in which ML-based AI systems have been most successfully applied have tended to be those that share certain characteristics with games. The scope of the task the AI is asked to carry out is limited and well defined, and large amounts of data are available to train the algorithms effectively. Applications ranging from facial and number plate recognition, to the machine's ability to recognise and decode the phonemes that make up human speech, or even the visual characteristics that define everyday objects all, to a greater or lesser extent, exhibit the same combination of characteristics.

However, in areas such as object recognition by autonomous vehicles, which must operate to extremely high levels of accuracy to meet safety concerns, the performance of ML systems still remains below what will be required. The task of correctly interpreting every object the system encounters in our highly complex everyday environment is neither limited nor well defined. Consequently, even the most advanced image recognition systems fall short [as we explore in the interview that follows with Dr Ali Shafti].

A chatbot can successfully handle simpler banking or insurance enquiries, because the range of tasks it needs to perform is limited by the nature of the conversation and the data it requires to do so is readily available from the customer or the bank's records.

But if a lonely bank customer were to phone the call centre in search of someone to talk to, a skilled and empathetic human operator could meet their need. Today's chatbots would stand no chance – the task is way beyond their ken. The best they could do would be to hand over the call to a human.

The Himalayan tasks to come

"Speech recognition is a good example where we have quite good performance now," says Prof Barber. "But it's still very superficial – the machine doesn't actually understand in a deep sense what you're saying. Similarly, translation appears largely successful. You can now get pretty good translations automatically from one language to another, particularly if the languages are quite close. But does it really understand what you're saying?

"There's nothing wrong with where we are right now. The progress we've made is extremely impressive but we're just at the foothills of this Himalayan task. We still don't know how to make that leap to systems that are much better at deep understanding. It's a leap that the tech giants are well aware of and are investing heavily to crack, because if you can crack that, then the utility of things like digital assistants will vastly increase."

Besides AI's inability to understand context, intuit or unspoken meaning, other challenges loom large. The data-hungry nature of ML-based systems necessarily limits the contexts in which they can be used most effectively. In data-poor environments, they struggle. This is one of the chief problems with so-called reinforcement learning, where AI-based systems learn from their environment rather than by crunching through huge volumes of training data learning to associate decisions with their long-term consequences. A man-made system that could learn from environmental stimuli, in the way humans do, would have to rely on a far smaller volume of information than current AI systems require. Making that leap to a more data-efficient style of learning is a key current research goal.

"Machine learners are intellectually fascinated by the idea of reinforcement learning because it is in some sense the 'mother problem' for AI: how to train systems with only very limited feedback on the eventual success or failure of a current decision," says Prof Barber.

The future of AI

Even though the prospect of "artificial general intelligence" remains distant, Prof Barber argues that real-world applications of AI, for example, in fully autonomous vehicles, are going to be hugely economically significant. Similarly, the development of robots that can accurately pack goods for despatch from warehouses, which remain largely manual environments, will have a major impact.

> "The progress we've made is extremely impressive but we're just at the foothills of this Himalayan task. We still don't know how to make that leap to systems that are much better at deep understanding. It's a leap that the tech giants are well aware of and are investing heavily to crack, because if you can crack that, then the utility of things like digital assistants will vastly increase."

Proponents of AI argue that widespread application of robotics and AI in business will free humans from many of the dull, repetitive, physically demanding tasks that we currently perform simply because human labour remains cheaper than robots. "Ever since the industrial revolution and even before that, we've been using humans as if they were machines," says Prof Barber. "Is it a good use of a human being to be stuck behind the wheel of a truck eight hours a day for 30 or 40 years? Humans are much better than that. We have amazing abilities – empathy, compassion, creativity. These are things that machines are rubbish at and probably will be rubbish at for a long time. I'm passionate about freeing humans to do the kinds of things that we are uniquely good at. So in that sense, AI is a very positive pursuit."

The transition he sees unfolding will involve humans working increasingly alongside AI-based machines, performing the tasks requiring greater experience and skill, while leaving the mundane, repetitive elements to the machines.

In areas that still involve largely manual processing of standardised tasks, such as the back-office operations of large banks and financial institutions, there is obvious potential over time to replace tens of thousands of human roles with so-called robotic process automation, a repeat of the automation of factory floors through the 20th century.

The prospect of an AI-induced transformation of the workplace provokes understandable fears, Prof Barber acknowledges. "I don't know if revolution is necessarily the right phrase. These things are always over-hyped. I think it's going to be an evolution to some extent and I think humans are always remarkably resilient in creating meaning in their work and personal lives, despite the transformations that happen in society. Not surprisingly, he warns against getting "too fearful" about the number of today's jobs that could be replaced by AI and robotics.

However, these issues will not disappear. The big questions for investors, companies and their employees will be less concerned with the potential effects of AI and robotics on economic activity, which in time will prove profound, but much more concerned with how far and how quickly governments and regulators choose to respond to the searching questions AI will pose for our societies.

What is 'true AI'?

As the buzz around this branch of computer science has intensified over the past few years, terms such as AI, along with related techniques such as ML, have been bandied around freely. AI is now often applied to projects that purists would not regard as AI, even though they use some of the same tools, such as the statistical analysis of very large data sets. Given the levels of excitement among businesses, governments and the public about the potential of AI, attempts to cash in on this wave of interest and excitement are not surprising. Indeed, some might argue on this basis that algorithmic trading does not constitute "true AI", but is simply data analysis on a huge scale.

In the academic community there is a clearer consensus about what constitutes "true AI", says Prof Barber. "AI for me is about the ability to replicate human perception and reasoning and our abilities to interact with each other and the physical world."

Does this blurring of boundaries matter? Prof Barber argues that for academics it is important to be clear about the scope of the discipline, not least so that funders can understand what their money is paying for. But beyond that, other issues weigh more heavily: "What's more important [than rigid definitions] is that we actually make progress in making systems that are of practical utility for mankind, things that people find interesting and really make a positive difference to our lives."

Professor David Barber biography

David Barber is director of the UCL Centre for Artificial Intelligence, which aims to develop next generation AI techniques.

He has broad research interests related to the application of probabilistic modelling and reasoning.

He is also chief scientific officer for re:infer, which is a natural language processing startup that "turns unstructured communications into structured data to drive action."

He received a BA in Mathematics from the University of Cambridge and subsequently a PhD in Theoretical Physics (Statistical Mechanics) from the University of Edinburgh.

To find out more visit

columbiathreadneedle.com

COLUMBIA THREADNEEDLE INVESTMENTS