New Leadership
for a *Changing* World

The Evolution of Artificial Intelligence

Paper One (of Two)

MILES ADVISORY

Welcome to this first part of a two-part series that looks at artificial intelligence (AI) and its implications for leadership and management. Other than 'Wordle', 'AI' was one of the most searched for terms in 2022 / 2023. Attention and interest in the subject have risen to a state of relative frenzy, particularly on business orientated social media sites like LinkedIn and Maimai (China's most trafficked professional social networking site).

This brief series of papers hopes to establish where we are with AI, and its implications for leadership and management in the short to medium term. We will try and clarify an extremely ill-defined and contested concept, outline some of the key principal components of AI, explore its existing manifestation in the world of business, and assess its implications for the role of leaders and managers in the AI enabled organisations of the future.

We will start this first paper with an overview of the history of AI, which helps to identify some of the key concepts we need to understand. This will necessitate the use of some relatively specific and technical terms. In our subsequent paper we will review how AI is currently manifesting in organisations and the implications of its evolution for practising managers and leaders.

A long history

Despite the overwhelming attention that AI has experienced in recent years, the core ideas have been with us for decades.

In Science Fiction, it has been a key topic for discussion and source of inspiration for well over 100 years; Alan Turing and John von Neuman laid its theoretical foundations in the 1940s and in the 1950s applications like the General Problem Solver (GPS) began to emerge alongside 'expert systems' - early attempts to allow machines and the programmes that ran them to 'learn' from experience and make better predictions based on gathered data.

By the early 1960s, computer scientists were focused on developing programmes allowing language comprehension, and by 1965, early attempts like Shakey and Eliza were engaged in primitive conversations with humans.

During this time the research on expert systems moved to more prosaic subjects as programmes were taught to play chess and checkers, and by the 1980s concern was emerging that these systems were increasingly able to outperform human subjects in these narrowly focused tasks.

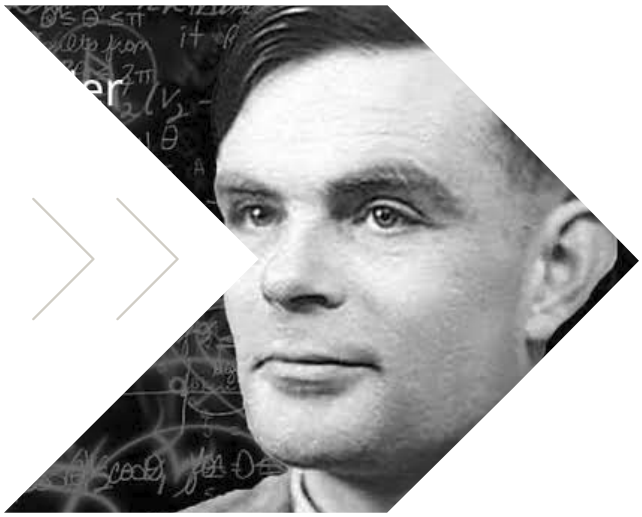

The 1990s saw huge advancements in computing power, information storage and the capabilities of computer hardware and software more generally. The cost of processing power fell exponentially (thanks to Moore's law), leading to renewed excitement and investment in the capabilities of what was increasingly being called artificial intelligence.

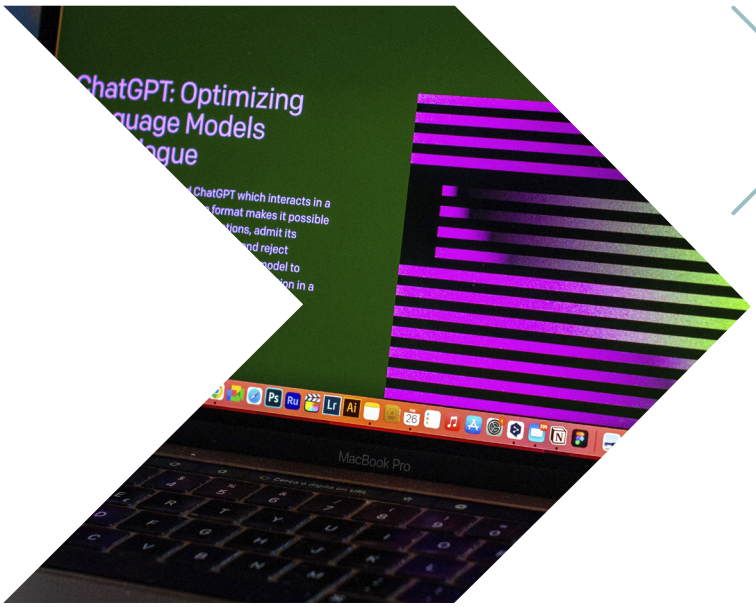

A critical development proved to be the interest of computer scientists in mimicking the working of the human brain to improve computing performance, leading to the first neural networks in early 2000s which took the idea of machine 'learning' to new levels.

Generative applications followed, based on reinforcement learning algorithms able to learn complex behaviour (like driving a car) based on relatively minimal data. As data became more abundant, 'training sets' used to instruct these programmes blossomed in size, breadth, and complexity.

In the last decade, executives in many fields of industry identified AI as the most potentially disruptive force in modern business (New Vantage Partners, 2016). In 2017, a study of senior executives around the globe by Accenture suggested that 85% of those executives intended to invest heavily in AI related technologies in the subsequent 3 years. Finally in November 2021, ChatGPT, a large language AI programme developed by OpenAI was released, and AI, in its most overt form, was finally in the hands of the population as whole, becoming the fastest growing consumer software of all time in January 2023, with over 100 million users worldwide.

Inside the black box

AI draws on many disciplines for its major theories and concepts including computer science, mathematics, psychology, linguistics, probability theory so perhaps it is inevitable then that the discipline is replete with often obscure and technical language.

Firstly, can we actually define AI? There are many different definitions, and debate rages about the most accurate and useful. Terms such as reactive, theory of the mind, limited memory, self-aware are regularly coined in attempts to structure such a definition. One suggestion is that AI is best described as a 'bad choice of words' coined in the 1950s (specifically at an academic workshop in Dartmouth in the USA held by John McCarthy, one of the founders of AI research). Generally, definitions seem to gravitate towards an attempt to simulate human intelligence and other key cognitive abilities in mechanical (or at least computational) terms. Hence AI attempts to encapsulate and incorporate functions such as learning, reasoning, problem solving, perception and language comprehension outside of the human brain.

Here we encounter our first key question and definition, how smart are we designing these systems to be?

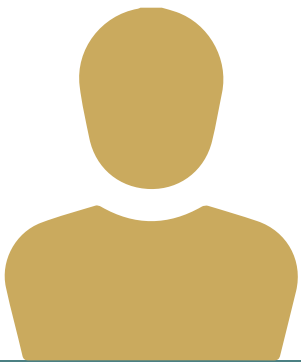

'Weak' AI systems (also called **ANI** – artificial narrow intelligence) typically have the capacity to 'learn', that is to take feedback and improve their performance, but within a narrow range of applications designed to perform a specific or very limited range of tasks (a logical extension of early 'expert' systems).. Siri, Alexa (etc.), recommendation algorithms and basic image recognition systems are typical examples. This is where we currently are with AI.

The next evolution of AI, according to this classification model, will be 'strong' AI systems, possessing 'human levels' of intelligence (and perhaps inevitably surpassing it), and able, through learning, to adapt to a range of different tasks. This idea (also called AGI – artificial general intelligence) is largely theoretical at the moment according to most researchers but may eventually evolve into its final form, **ASI** – artificial super intelligence, self-aware systems, capable of a huge range of tasks, but essentially the stuff of science fiction and fantasy today.

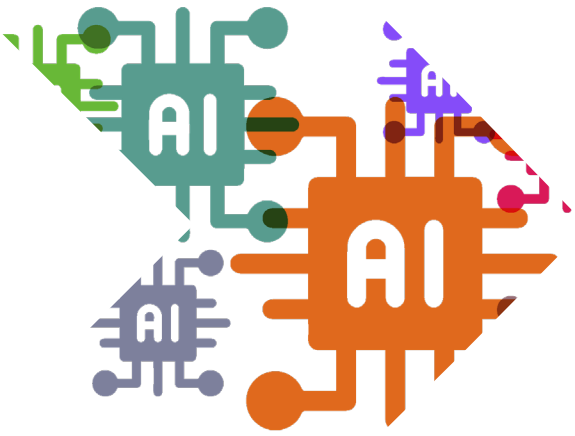

We have used the term 'learning' extensively so far, but what does this mean in terms of AI? Machine learning was first used as a term in the 1950s. It essentially describes the use of algorithms designed to allow programmes to take feedback on their outputs and adapt themselves to improve their performance. Here we hit another widely used term, algorithm. What does it mean? In essence the idea is very simple, an algorithm is a set of instructions, that is step by step procedures to solve specific problems – a finite set of instructions carried out in a specific order by a machine. Whist simple in concept, there are huge variations in different types of algorithms used in computing - in any discussion on their relative merits you are likely to hear terms such as brute force, recursive, backtracking, searching, and sorting. Applied to the specifics of AI, and particularly machine learning, the terms will expand to cover decision trees, random forests, boosting, support vector machines, k-nearest neighbours etc.

Within the scope of machine learning in general, we have 'Deep' learning, typically based on neural networks. Deep learning focusses on training artificial neural networks inspired by human brain structure and functioning.

What is a neural network (or ANN – artificial neural network)?

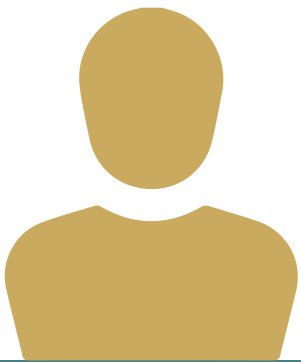

Neural networks emerged from attempts to improve the performance of computer programmes by studying the workings of the human brain. Drawing on biological phenomena such as neurons, designers utilised artificial neurons or 'nodes', signalling paths connecting different data working together to solve problems.
A typical neural network will have at least three levels of nodes: input, hidden and output (the 'deep' part of deep learning refers to the depth or number of layers of nodes). Nodes have weights and thresholds, and much like the human brain, trigger each other into action. For those with a more mathematical bent, think of each node as a linear regression model. The network of nodes relies on training to improve their performance over time. A relatively early example is Google's original search algorithm.

Critically, when compared to traditional machine learning models, which tend to plateau in performance over time no matter how much data is thrown at them, neural networks continue to improve in terms of performance as they are fed more data.

They can automatically extract relevant features from raw data, drawing on multiple interconnected nodes in the network and have proven to be particularly useful in computer vision and natural language applications. Hence ChatGPT is described as a feed-forward 2 neural network which has an additional layer of nodes, focused on 'normalization', interpreting outputs to form human-like responses.

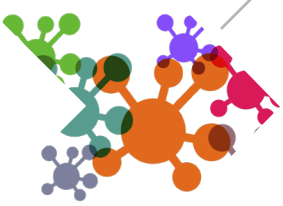

Some of the more recent variants of AI to emerge include:

GANS – Generative Adversarial Networks. The term **generative** means simply that, programmes that can generate content such as text, images, music, audio and videos.

DALL.E (focused on image recognition and manipulation) and natural language applications (like **ChatGPT**) rely on **'Transformers'** (the T in ChatGPT) allowing programmes to focus on elements of language and images to mirror 'real life' more closely. They are 'adversarial' as two AI programmes, a generator and discriminator, work together (essentially compete with each other) to produce the most life like image or language responses. Such programmes are typically those used to devise the 'deep fake' images which are currently sweeping the internet.

Other recent variants include ethical AI and emotional / sentiment AI, though both are in relatively early stages of development.

Typically, for these networks to learn, they need to be supplied with data. Up until relatively recently, the data has often been fed to such programmes by a user, labelled appropriately, typically referred to as 'supervised' training.

Such 'training sets' have increased hugely in size and complexity over time, to the extent that it may now be easier to allow the AI programme itself to collect its own data – unsupervised training. Here we enter a debate that troubles many commentators on AI – what data are these applications using?

Data availability and quality are clearly critical. Bias (conscience or unconscious), improper labelling, privacy issues can all affect the quality of data inputs and hence programme outputs, particularly in the case of neural networks. To what extent are we building in our own cognitive limits and biases into the data and hence the applications themselves?

If such programmes are allowed to generate their own sources of data, in an attempt to minimise the apparent 'dangers' of human interaction, what data will they be using to generate their outputs?

Perhaps ultimately, this points to what should be our biggest concern about AI, that it is being developed and used by us, humans, with all of our unconscious biases, good and bad intentions, overwhelming egos, and the frailties and imperfections that make us human.

The most significant question we all face may be are we up to the job of guiding the birth and maturation of such a potentially significant and transformative technology in our modern society? These applications will remain 'black boxes' for the vast majority of users, hence we have only the dimmest idea of how they operate, the data they are working on and how they are manipulating it to generate outputs.

The literature surrounding AI is, as we have seen, replete with terms such as intelligence, self-awareness, consciousness etc. It is worth noting, that even after hundreds of years of research, we still don't have generally agreed upon definitions of these concepts in humans and other biological species, let alone machines.

We still don't understand the 'black box' of the human brain. Many commentators are concerned at the prospect of these applications becoming self- aware, conscious, and working to their own agenda, rather than ours, but will we know when this happens? We don't really know how to measure it.

Creating the appropriate environment in which AI can evolve in the modern world and contribute to social welfare (as well as profits) is potentially our biggest challenge in the short term. In the next paper we will look at what leaders and managers can do to guide this process in their own organisations.

References in this article

Moore's law – the number of transistors on computer chips doubles every year, leading to significant improvements in computer processing power

That is information only moves in one direction – forward – through input, hidden and output nodes.

What next?

We would welcome an opportunity to meet with you and engage further and believe in building relationships over the long haul.

Chris Stainton - Partner

M: 00 44 7957 434736
E: chris.stainton@miles-advisory.com
20 North Audley Street
Mayfair
London
W1K 6WE

www.miles-advisory.com

linkedin.com/in/chris-stainton-2b983311

Chris is a Partner and has a wealth of experience leading senior search assignments across infrastructure, consumer services, life sciences and technology sectors both in the UK and internationally.

Additionally, Chris has advised clients on Board transition, effectiveness, and a broad range of leadership solutions.

Clients include both publicly listed companies and private equity. Chris began his career in recruitment in the consumer goods sector before moving into Executive Search in 1998.

MILES ADVISORY

disability confident EMPLOYER