Chapter 2

Artificial Intelligence

Artificial Intelligence

Key points

- Artificial intelligence (AI) is the technology that helps a machine to think like a person.
- Artificial intelligence is already being deployed within banks, to interface with consumers and businesses, and within government itself.
- Machine learning—a particular type of AI—has offered new opportunities in financial services and shaped new offerings.
- Inevitably, however, with new opportunity comes new risks, and regulators need to address these proactively.

2.1 Introduction

Commercially scalable artificial intelligence (AI) has become widespread over the past decade. The financial services sector adopted it quickly, and it has since created both opportunities and risks that central banks need to consider. There are several implications for the use of AI within a central bank and for the adoption of AI by financial services institutions and in other industries.

2.2 Context

Artificial intelligence is a term we use to describe machines that can think in highly sophisticated ways—ways that aim to imitate (or surpass) human cognition. Some notable examples of AI are Google's DeepMind, IBM's Watson, Amazon's Echo and Apple's Siri.

While AI is the superset, subsets of AI that we might hear discussed in relation to financial services include applied AI (AAI), generalised AI (GAI), deep learning and machine learning. We will look closely at the last of these in the next section.

From a regulatory standpoint, it is helpful to think of major deployments of AI within:

- a supervised bank or non-bank financial institution (NBFI);
- the market (whether consumer- or business-facing); and
- government, including the central bank.

We will look at these in brief, before considering the future of AI in digital financial services (DFS).

2.3 Description

2.3.1 Artificial Intelligence, Machine Learning and Deep Learning

Machine learning is a subset of AI that involves supplying machines with enough data that they can learn what it means or learn to interpret it. This contrasts with traditional computer programming, which involves giving a computer a discrete set of instructions that it simply follows (a 'rules-based system'). The terms 'AI' and 'machine learning' can be confusing because they are

often (incorrectly) used interchangeably. When banks or entrepreneurs, for example, talk about AI, they are usually talking about machine learning.

One example of machine learning is the natural language processing that powers applications such as Apple's Siri, Amazon's Echo, Google Home or Tesla's Autopilot. By integrating vast quantities of data about what people say and how they say it, the devices are able to interpret meaning and, for example, order a pizza when we say the words 'Order me a pizza'.

By contrast, an example of a rules-based system is the system underpinning self-driving vehicles.

Machine learning is something that we already encounter in a number of applications today. Generalised AI, meanwhile, is still several years away, although great strides have been made. DeepMind, for example, grew out of a research effort at University College London's Computational Neuroscience Unit. The Unit's goal was to 'solve intelligence' and the team approached its research by combining machine learning with systems neuroscience (a discipline that involves understanding how neurons form networks). Unlike many other 'guided' AI systems, DeepMind is said to learn only from experience and not to follow predefined paths, relying instead on general-purpose learning algorithms. DeepMind and other efforts such as Watson are starting to find their way towards broader applications.

Figure 2.1 explains the relationships and differences between deep learning, machine learning and AI, and the evolution over time of these different systems.

As Figure 2.1 explains:

- AI is the broadest possible category, encompassing all efforts to make a machine imitate human thinking;
- machine learning is a subset, focused on computers that can learn to imitate human thinking without being programmed specifically to do so

Figure 2.1 The relationship between AI, machine learning and deep learning.

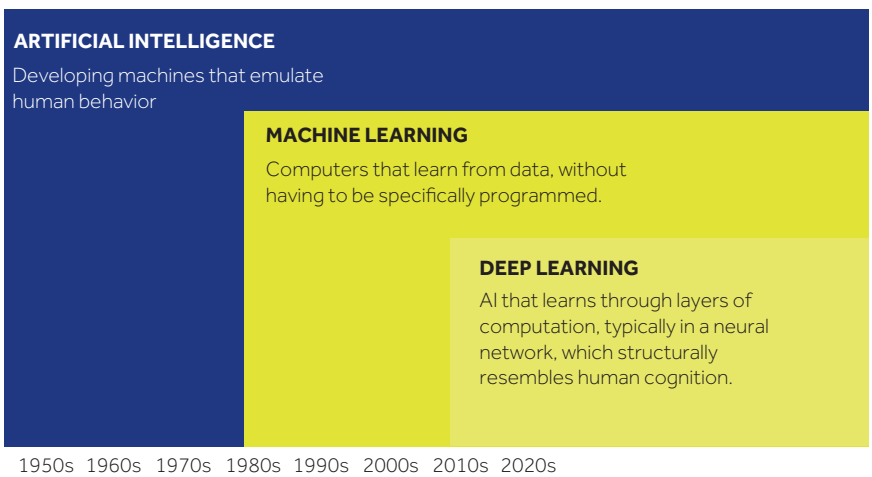

Source: D Shrier (2021). 'The Convergence Revolution'. In D Shrier and A Pentland (eds). *Global Fintech.* Cambridge, MA, and London: MIT Press. Reproduced with permission of the author.

(rather like how a human child learns language); and

- deep learning is a special kind of machine learning that can address more sophisticated types of problem and its work is much more accurate if it has access to sufficiently large quantities of data.[1]

2.3.2 How is AI (Primarily Machine Learning) Being Applied in Financial Services?

We can examine the main applications of AI in financial services and how those applications relate to the concerns of regulators by considering them in relation to the areas in which they are deployed.

Within a Supervised Bank or NBFI

Artificial intelligence—particularly machine learning—is streamlining financial services infrastructures. Banks and NBFIs can use it to assess risk, to monitor transaction flow, to implement cybersecurity and even to make critical decisions rapidly. Inevitably, however, each of these new applications involves potential new risks. Since, historically, many AI systems have operated as 'black boxes' within which where and how they make decisions is obscure to the average human being, it can be difficult to audit those decisions and attest to their compliance.

Some financial institutions have begun to use algorithmic models to make core lending decisions such as credit underwriting. If applied with care, this type of model creates an opportunity for the 3.5 billion people who are underbanked or unbanked because machine learning can help to break the 'credit trap'—that is, the fact that an individual or business has to demonstrate a credit history to get credit in the first place. New models based on human behaviour and alternative datasets are now making it possible for a bank to extend credit responsibly to first-time borrowers.

If not properly trained or supervised, however, machine learning systems can end up configured such that they further exclude already-marginalised groups on the basis of characteristics such as race or low incomes, so deepening the divide that they have the potential to bridge. To militate against this risk, regulators may require banks to submit to an annual algorithm audit, certified by a third-party firm, which should examine particular areas of compliance and risk.

Within the Market (Consumer- or Business-facing)

Artificial intelligence systems offer numerous opportunities for inclusive consumer and business financial services. For example, in Bangladesh, an effort is under way to implement voice-enabled technologies, opening up access to financial services to that portion of the population whose literacy is poor. Teaching the machine learning system to recognise numerous local dialects and accents therefore means training it in a local context. A speech recognition engine trained only in received British English pronunciation, for example, would fail to meet local needs.

Large tech companies may lay digital identification and biometrics over AI pattern recognition systems, and this can create issues in implementation. For example, Google initially trained its facial recognition algorithms—an expression of machine learning—largely on pictures of people of European descent. When confronted with a photo of a person of colour, the systems produced errors that did reputational damage to the company—and unfortunately the steps it took to remedy the issue were likewise controversial.

When deciding on policy that will guide AI applications in DFS, it is therefore essential that regulators understand both the principles and the implementation of such systems.

Within Government, Including the Central Bank

The technology that financial services regulators use to manage processes such as monitoring compliance is known as regtech. Regtech has been lagging behind the fintech industry in several ways, but efforts are now being made to address this. Artificial intelligence systems could help government regulators and policy-makers to collect revenue, to monitor the financial system, to highlight or predict risk events, and in other ways to more readily accomplish their mission.

One area in which governments have been successful is in using large-scale datasets (including financial data) to develop and implement government policy through efforts such as Data for Development and, more recently, the Global Partnership for Sustainable Development Data (GPSDD). These initiatives saw government ministries assemble and sandbox large-scale datasets. They then invited hundreds of academic research groups from around the world to develop insights into issues ranging from inclusion to public health based on these datasets.

2.4 Key Considerations for the Future

It is increasingly important to align the use of AI with the values and cultural norms of any given country. Certain intrinsic bias is built into any such systems, so it is important to understand what this bias is and how it intersects with legislative and regulatory priorities. For example, all credit underwriting contains known bias that the regulator has deemed acceptable.

> "Artificial intelligence systems could help government regulators and policy-makers to collect revenue, to monitor the financial system, to highlight or predict risk events, and ... to more readily accomplish their mission."

Similarly, newer AI systems (whether for credit or for other applications) will require a regulator to understand how the algorithm arrives at its decisions if that regulator is to ensure that the application complies with relevant law and regulation. In fact, this form of 'explainable AI' is a growing subset of the broader AI commercial world.

Central banks and regulators will want to increase the sophistication with which they approach AI. They should appoint working groups to ensure that they stay up to date with developments in the field, and they should ideally train and employ in-house experts who know how to evaluate

and monitor AI systems. There are also opportunities for governments to support economic development and to stimulate or focus industry activity in particular areas to encourage private industry to apply AI towards desired policy outcomes.

Endnote

1 Mahapatra S (2018). 'Why Deep Learning over Traditional Machine Learning?'. *Towards Data Science*, 21 March [online]. Retrieved from: https://towardsdatascience.com/why-deep-learning-is-needed-over-traditional-machine-learning-1b6a99177063