Rainbird

AI in context

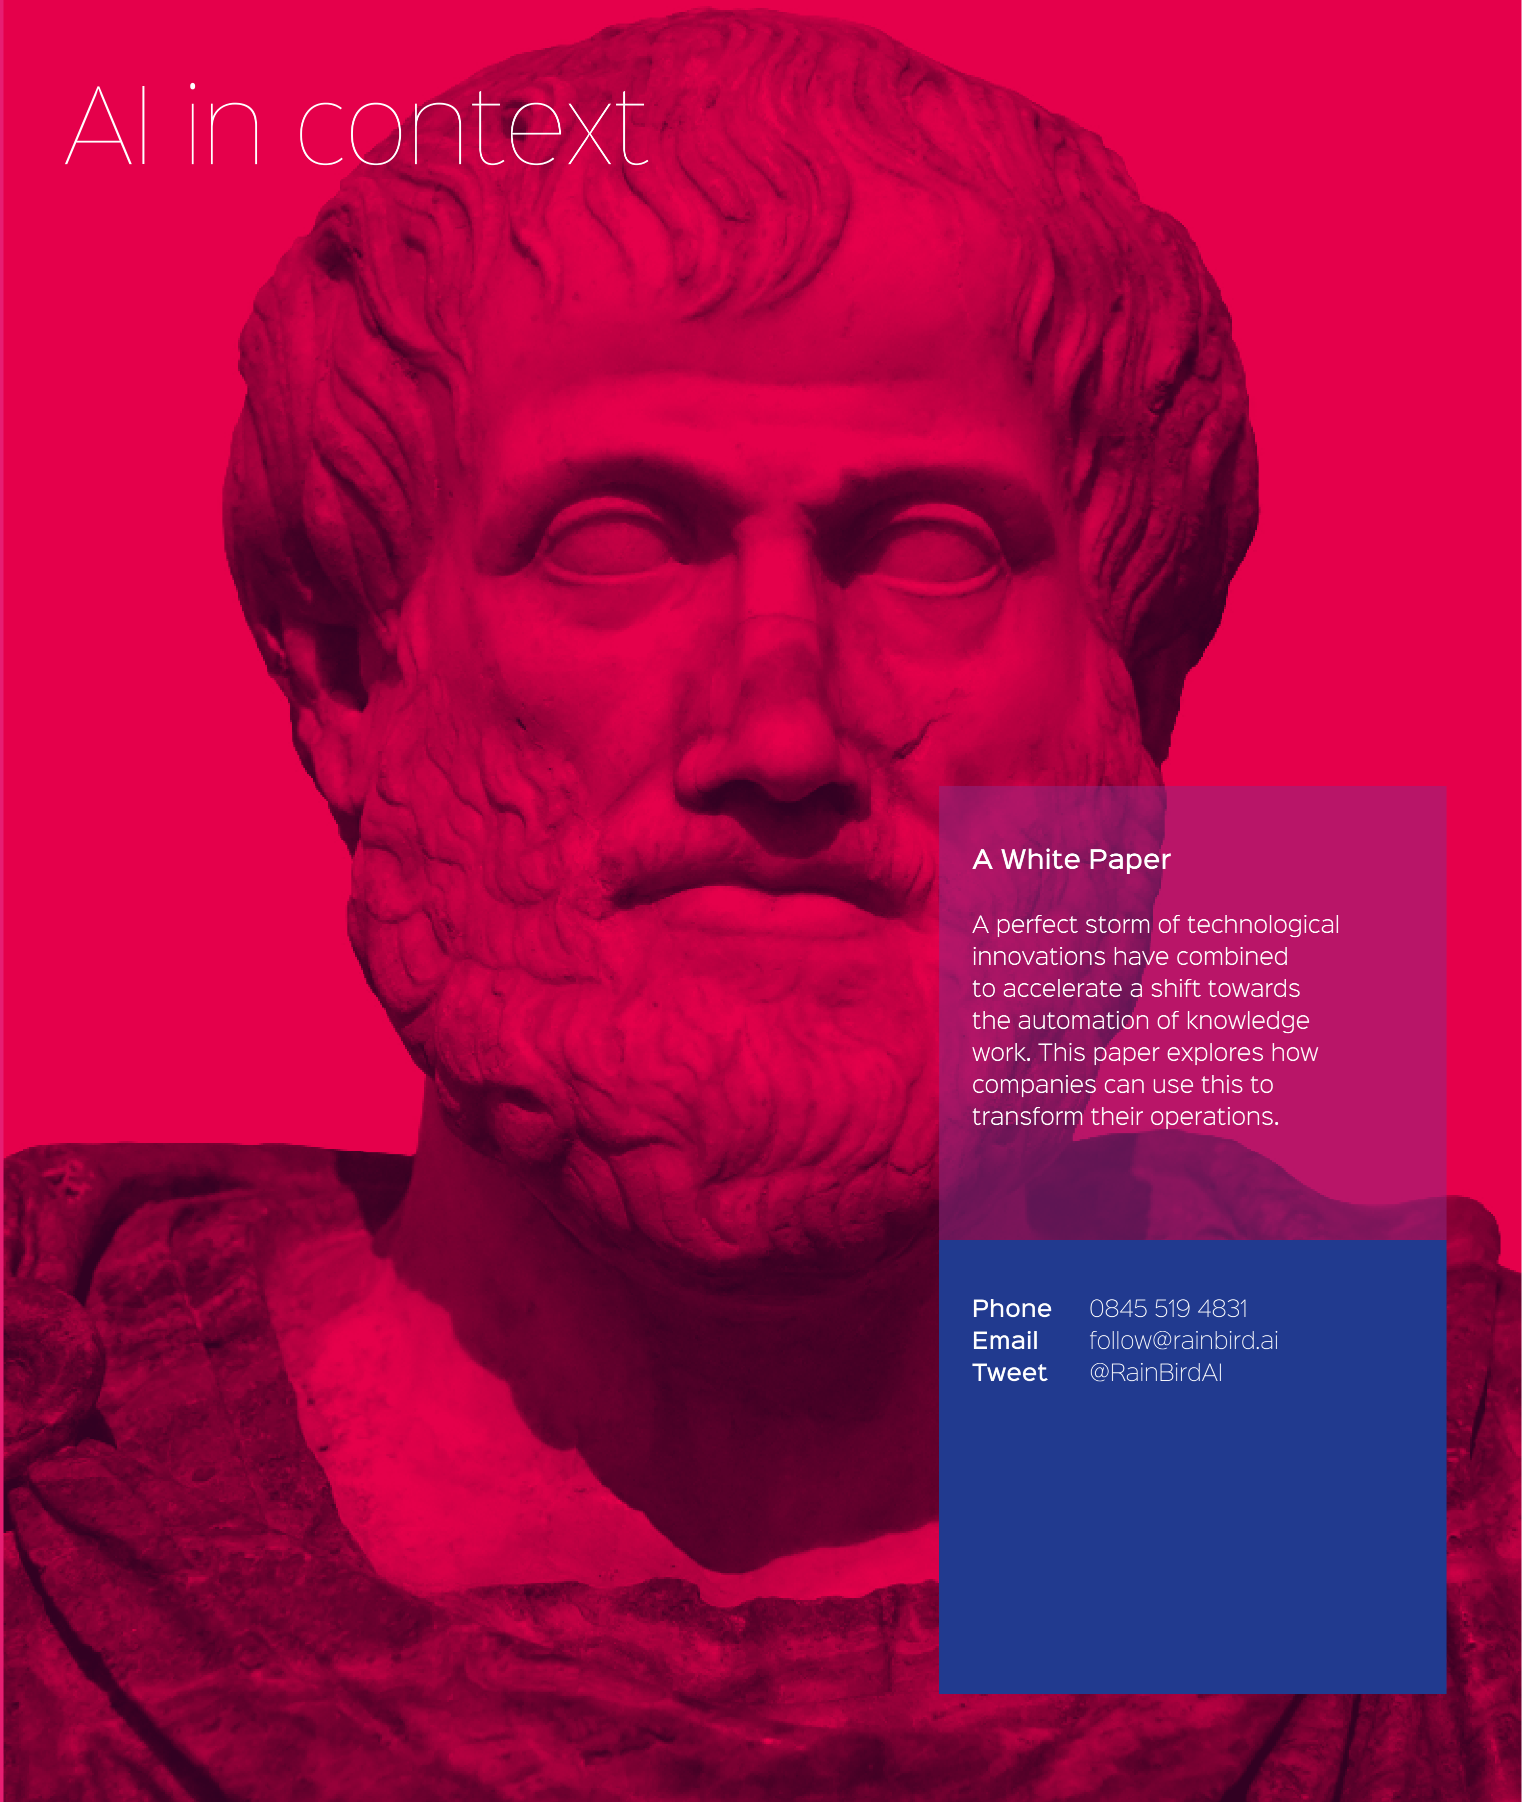

A White Paper

A perfect storm of technological innovations have combined to accelerate a shift towards the automation of knowledge work. This paper explores how companies can use this to transform their operations.

Phone 0845 519 4831
Email follow@rainbird.ai
Tweet @RainBirdAI

AI in context

Artificial Intelligence (AI) has seen something of a resurgence in the media in recent years. Despite this, the field remains poorly understood and is shrouded in mystery and obtuse terms. It can be hard to see how organisations can practically apply the recent advancements in AI.

AI is quickly becoming one of the most important technologies ever invented. Organisations that can understand the implications and act on them stand to gain significantly. This paper provides an overview of this rapidly evolving field.

We'll look at the history, current state and likely futures of AI and touch briefly on the debate surrounding the ethics and societal impacts of the technology.

Early history

It's easy to think of AI as a recent innovation, that this field of research came with the advent of electronic computers. In fact AI has it's roots much further back, with some of the early philosophical thinkers.

In around 350BC Aristotle devised syllogistic logic, the first formal deductive reasoning system. This system is a form of logical argument and uses a written representation of human-like deduction to reach a conclusion.

Aristotle and the many philosophers and thinkers that followed him, were concerned with finding a formal way to represent the way humans reason and think about the world. Indeed that could pass as a good description of one of the aims of AI.

In the early 20th century Bertrand Russell and Alfred Whitehead published Principia Mathematica, a work which laid down the foundations for a formal representation of mathematics. Given that formal representation, Alan Turing was later able to show that any form of mathematical reasoning could be processed by a machine. This was hugely significant at the time and led to speculation of the possibility of building thinking machines.

Complex brains

We know that the human brain is the most complex object in the known universe, and ironically is also the only object we know that might be capable of understanding the universe.

All of humanities' endeavours – agriculture, medicine, science, ethics, and space flight – have been the result of human cognitive function. One of the objectives of artificial intelligence research has long been to replicate or even improve on human cognitive function. Given that everything humanity has achieved so far is a result of that function, it's not hard to see the potential of AI.

However, no brain on earth is yet close to understanding how we might achieve this, despite all those centuries of trying. It's little wonder that the field of AI, the challenges it presents and the rewards it might deliver, have such a hold over us.

Electronic machines

From the dawn of electronic computing, the desire to build computers that can out-perform us at human tasks has been prevalent. This is another good definition of one of the goals of AI.

In the early 1950's Alan Turing wrote a paper titled 'Computing Machinery and Intelligence', which opened with the question 'Can a machine think?', and in which he proposed his famous test. The essence of the test argues that when a human can't differentiate an interaction with a computer from another human, that computer is intelligent.

The usefulness of his test is still being hotly debated six decades later, which points to how little we really understand the nature of intelligence. Despite the media attention, comparatively little research has been applied to solving the Turing test. Research in this area has instead been focused on understanding the underlying principles at work.

AI research accelerated in pace during the 1950's and 60's. By 1967, Marvin Minsky at MIT was predicting that, "within a generation… the problem of creating 'artificial intelligence' will substantially be solved.".

It seemed that there would be no stopping the advancement of AI, but Marvin Minsky would probably agree that, within his definition, we're still not there.

So what went wrong?

By the 1970s it started to become clear that the size of the AI problem was monumental. Limited computing power at the time combined with a reduction of funding as it became clear that AI research was failing to live up to the hype. This led to the first "AI winter".

Attempts to build systems based on the first order logic described by those early philosophers failed because they couldn't adequately describe uncertainty or ambiguity. Technologies we now recognise as the foundation of modern day AI research were, at that time, still being dismissed.

Modern AI

Despite this, the research didn't stop entirely. The development of expert systems started to achieve significant commercial interest in the 1980s. These systems leaned heavily on the thinking of those early philosophers and were centred on structuring knowledge.

A connectionist approach was also revived with neural networks starting to find commercial application in areas such as computer vision.

AI started to be applied to fields as diverse as stock trading, oil prospecting, agriculture and medicine.

In the case of medicine, expert systems were developed to support clinical diagnosis. Sample analysis, which had previously been a human intensive task, started to become supported by computer vision classifiers. Financial trading systems started using more and more AI as computerised systems became able to make decisions quicker and more consistently than humans. Automatic production planners and schedulers found widespread use in industry, hugely increasing the efficiency of manufacturing.

In 1997, IBM's Deep Blue became the first computer to beat a world champion chess player. This was seen as a major milestone, not least because being good at chess was one of the early promises of AI. It achieved this feat by sheer weight of computing power. Deep Blue was not the first chess playing computer, but now IBM had the computing grunt required to better humansat this single task.

However, building systems that made use of this technology was extremely costly. The advances that were made remained accessible only to the largest corporations or constrained to academic research. The development of expert systems still required a large team of specialist software and knowledge engineers to build and maintain. The cost of implementation frequently ran into millions of dollars.

Despite the costs, there were many successes which were in well targeted, specific and mostly commercial applications.

The AI Effect

Most of the commercial successes of AI in the 1980s and 90s are systems we would now recognise as mainstream computer science, and not AI at all. We would not now argue that Deep Blue was intelligent, but we would accept that it performed a task that was previously thought to be specifically human.

We can now accept that the best chess players are computers and no longer consider chess as a task that's specifically human. As a result, chess playing computers no longer fall within one of our definitions of AI: "to build computers that out perform humans at some task that was previously thought of as being specifically human".

Indeed John McCarthy, who coined the term AI in 1956, noticed this phenomena. He once complained that "As soon as it works, no one calls it AI anymore". This is generally referred to as "The AI Effect".

Because of this it's easy to fall into the trap of thinking about AI only in terms of science fiction or a relic from computing history that never lived up to the hype.

Clarifying AI

Two very useful new terms have been coined in order to differentiate efforts to build AI for a specific purpose from an AI that can be generally cognitive.

Artificial General Intelligence

Artificial General Intelligence (AGI), sometimes called Strong AI, is a research field attempting to build machines that are generally cognitive. They are not being built for a specific purpose, but are designed to perform the full range of human-like cognitive functions. Some argue that an aim of AGI is to build a conscious machine.

Much of this research currently focuses on building systems that replicate neurological processes.

Great strides have been made in recent years:

- The Human Brain Project is looking at the feasibility of representing the entire human brain.
- Google DeepMind can demonstrate a single machine that can learn to play thousands of different computer games.
- Advancements in specialised neuromorphic hardware can more closely mimic the way the brain works.

AGI also raises the question, "Is building an Artificial General Intelligence simply a matter of building a machine that models every connection in the brain?". If we achieved that, would we have a machine that behaved like a person? There remains a significant gap in our understanding that we must bridge before that question can be answered.

The debate that's currently raging on the ethics of AI research argues that we must regulate AGI to prevent us stumbling into building a machine that might eventually replace us. The argument states that, once we build a machine with human level general intelligence, it's likely it will rapidly surpass us. The fear is that AGI may become the last invention we ever make. There is much speculation over the implications of achieving AGI. The term "The Singularity" has been coined to describe what that future might look like because it is hard for us to predict a world where machines are more intelligent than us.

However, there remains a huge gulf between modern computing and human-level artificial intelligence. The research community remains extremely divided over when this might be achieved. Estimates range from the next 20 years to several centuries, with most researchers predicting long timescales.
While Moore's Law now looks to be slowing down, there are still many advances to be made in electronic processing and with optical and quantum computers now also in development. It seems certain that computing power will continue to increase to support the demands of the AI engineers. AGI remains a very active field of academic and commercial research with spin-offs finding their way in to all areas of computer science.

Artificial Narrow Intelligence

Artificial Narrow Intelligence (ANI), sometimes called Applied AI, is a term used to describe an AI that performs a single function as competently, or even better than a human.

IBM Deep Blue was an ANI that beat Garry Kasparov by being very good at searching for an optimum chess winning strategy. IBM Watson won Jeopardy by having access to unprecedented amounts of existing data. ANI is sometimes referred to as weak AI, a term which is unhelpful and less descriptive. We would not generally refer to either of these IBM machines as weak, just narrow in their scope.

There's no doubt these systems are outstanding at what they do, but neither could play Monopoly for instance. Certainly a system could be built that is very good a playing Monopoly, but that would be another very specialised machine.

The adoption of probabilistic representations and statistical learning methods has led to cross fertilisation between previously disparate fields. AI, statistics, control theory, and neuroscience are all coming together in a perfect storm of innovation.

What's more, society is starting to accept these narrow applications of AI and it's already becoming normal to devolve more and more of our thinking to machines. Few people like repetitive, rote tasks and adopting technology frees the workforce to perform more meaningful work.

Many of the problems that AI suffered in the 1970s can now be solved by building narrow applications of AI. We can overcome the shortfalls of the logical approach to knowledge representation by using probabilistic programming techniques. Many of these early methods also failed because data sets were tiny and computing power was millions of times too slow by comparison with today.

ANIs are now changing the world and the impact to business has already been huge:

- More than half of the equity shares traded on US markets are traded by algorithmic high-frequency AI traders, not people.
- Autonomous self-driving cars are now a reality and being tested on British roads.
- Siri combines several ANI technologies enabling our voice commands to be decoded, interpreted and acted on within narrow parameters.
- Satellite navigation systems have long since been able to plot an efficient route across a city based on traffic analysis.

These are all examples of ANI.

Knowledge Work Automation

If we look back over the last century, we see the significant automation of physical work as machines started taking on heavy-lifting and repetitive tasks. This automation started in agriculture and quickly spread to manufacturing and other sectors.

Automation has continued to accelerate and is spreading to more knowledge-based jobs.

In the same way that machines have changed the way we have approached physical work, this new generation of AI-powered technologies are increasingly to be found in the workplace aiding the knowledge worker. Knowledge work automation is rapidly changing the workplace.

As with previous technological innovations, ANI and computerisation represents a step change in the nature of the work performed by humans.

McKinsey predict that over the next decade the impact to the global economy of automating knowledge work will be more than $5-$7 trillion per annum. That's equivalent to 110-140 million jobs.

We are just at the beginning of this tidal wave of change. Computing power is increasing and continues to get cheaper, and very big data is becoming more accessible and available.

As computing power has increased, and our computers have become more connected, so too has the accessibility of these technologies. It's no longer necessary to employ large teams of software developers to build systems that can radically change the way you approach your business.

Rainbird is an example of a technology that enables businesses to create their own Artificial Narrow Intelligence solutions.

To learn more about how Artificial Narrow Intelligence solutions can benefit business, request our White Paper The Automation of Knowledge Work.

History of AI

- **350BC** – Aristotle devises Syllogistic Logic
- **1910** – Bertrand Russell and Alfred Whitehead publish Principia Mathematica
- **1936** – Alan Turing proves any mathematical reasoning could be processed by a machine
- **1950** – Alan Turing writes Computing Machinery and Intelligence – Turing Test is devised
- **1956** – John McCarthy coins the term ‘AI’
- **1970** – First ‘AI Winter’
- **1980** – Expert systems achieve commercial interest
- **1997** – IBM’s Deep Blue becomes the first computer to beat world champion chess player
- **2011** – IBM’s Watson wins Jeopardy
- **2013** – Rainbird founded
- **2014** – Google release self-driving cars

Phone 0845 519 4831
Email follow@rainbird.ai

www.rainbird.ai