Cambridge Judge Business School

CONSUMER ARTIFICIAL INTELLIGENCE

The unintended consequences of implementing artificial intelligence for personal use

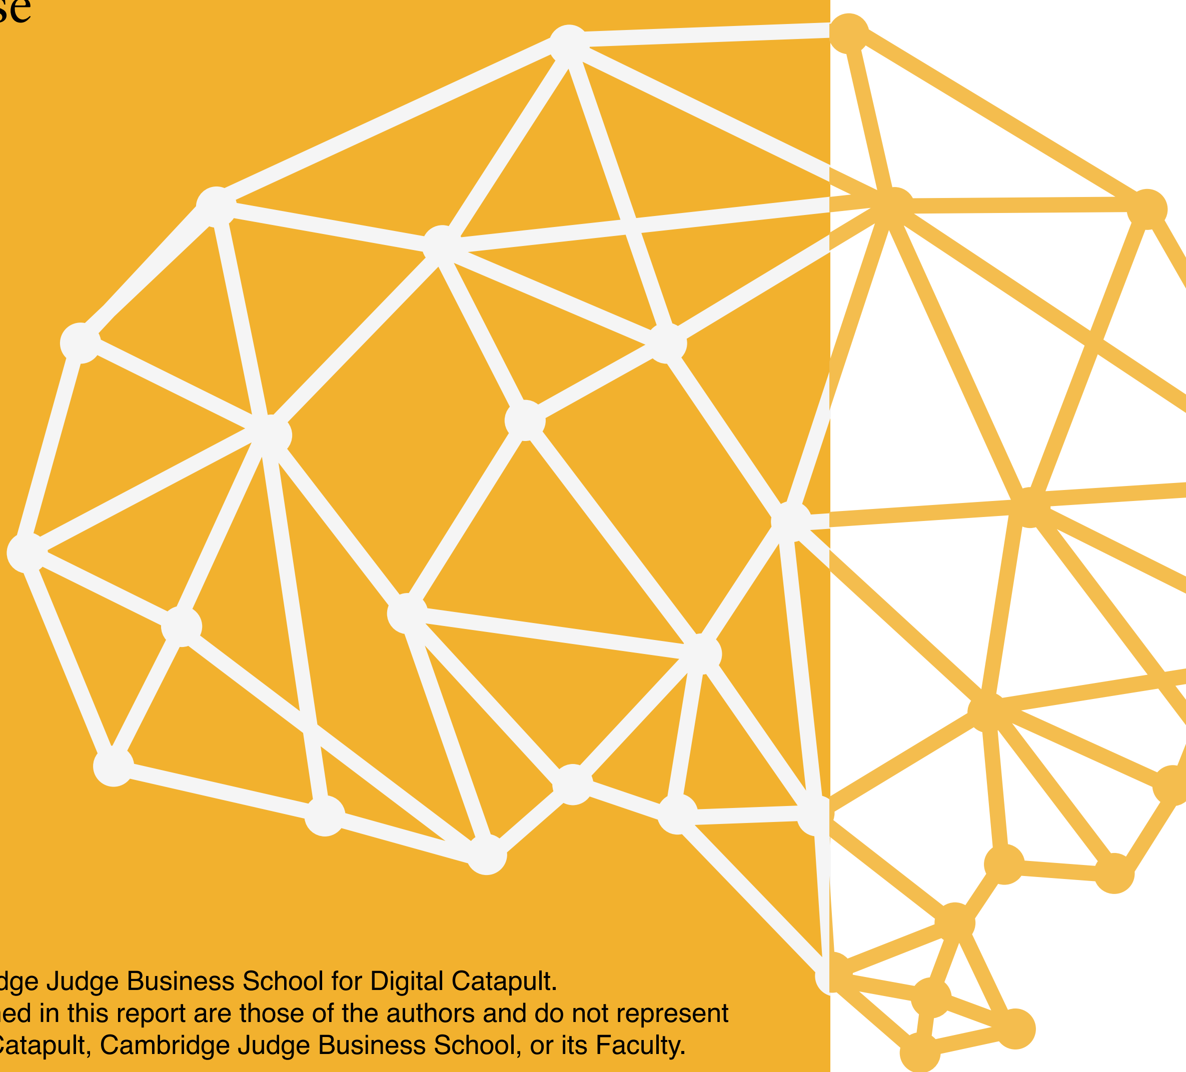

May 2017

Prepared by Cambridge Judge Business School for Digital Catapult.

UNIVERSITY OF CAMBRIDGE
Judge Business School

CATAPULT Digital

Authors

Aleksandra Pedraszewska

Hadi Michael

Mohamed Foulser

Satenig Mirzoyan

Seth Cleary

Zachary Dannelly

Supervised by

Vincent Mak
Reader in Marketing & Decision Sciences, Cambridge Judge Business School

Brian MacAulay
Lead Economist, Digital Catapult

Anat Elhalal
Head of Technology - Artificial Intelligence and Machine Learning, Digital Catapult

Phil Young
Senior Policy, Research & Engagement Manager, Digital Catapult

Acknowledgements

UNIVERSITY OF CAMBRIDGE Judge Business School

CATAPULT Digital

With thanks to individuals from the following organisations

Ai BUILD

BIOBEATS

blend

COM PLEX

Cyberlytic

esa

Flexciton

IBM Watson

intel

CFI LEVERHULME CENTRE FOR THE FUTURE OF INTELLIGENCE

Mind the Bridge

NORTH DAKOTA UNIVERSITY SYSTEM

Oxford Internet Institute University of Oxford

PicsArt make awesome pictures

ripjar

skim.it

SOUNDCLOUD

The Alan Turing Institute

UBER

Cabinet Office

Department for Business, Energy & Industrial Strategy

Department for Culture Media & Sport

UK Representation to the EU Brussels

United States Naval Academy Annapolis

zendesk

TABLE OF CONTENTS

Executive Summary

The field of artificial intelligence (AI) has been attracting unprecedented interest from business, governments, and academia (Lunden, 2016; Simonite, 2017; Smith, 2016; Agrawal et al., 2016; The Economist, 2016d). This trend has been fuelled primarily by: the emergence of Big Data, extensive research in Information Technology, and ever-growing compute power (Kelly, 2014). It is no surprise that, as a result, AI applications are no longer restricted to computer science labs but find their way into marketplaces. Despite that, there is no consistent source that businesses and policymakers can refer to in order to build their understanding of the nature and effects of these consumer-facing AI applications. This report fills this gap by focusing on the commercial developments of intelligent agents for personal use, defined by the authors as Consumer AI.

Based on comprehensive primary research, this original analysis is structured around the notion of unintended consequences. Thus, it focuses specifically on the benefits and drawbacks beyond the intrinsic motivations of Consumer AI implementers. Key findings were derived from a combination of semi-structured interviews with 30 experts across business, government, and academia regarding the topic of AI, and a survey of 300 consumers based in the United Kingdom. Secondary sources and case studies identified in the literature provided context and further evidence. As a result, this report identifies ten interest areas where unintended consequences of Consumer AI are evident. As such, it contributes to the literature by identifying the interest areas in which the consequences are most applicable to Consumer AI, and maps the effect on consumers as direct or indirect.

This report (1) shows evidence of overconfidence in AI consumers' knowledge, underlying an opportunity for both the private and public sectors to provide more information on the workings of AI-reliant products and services. We present a cutting-edge application of (2) the traditional principal-agent problem with respect to the relationship between consumers and AI agents, which may lead to adverse selection and misalignment of commercial and consumer interests. This report shows that the questions of agency extend beyond this relationship. The increased use of some AI techniques in the Consumer AI realm results in (3) the agent scrutiny problem, and thus requires a distinct regulatory approach - on agents' inputs and actions rather than their function or program. Next, we consider (4) discrimination resulting from prejudice in the development and training of intelligent agents, (5) the heightened impact of data breaches, compromised agents, and offensive capabilities requiring more severe cyber-security measures, and (6) the potential for Consumer AI to affect market structure, leading to anti-competitive behaviour.

Other key interest areas in which unintended consequences of Consumer AI are evident, while still applicable to the broader field of AI, include: (7) a shift of academics to the private sector and hence privatisation of future AI research, warranting an increase in governmental funding, (8) the democratisation of the AI developer ecosystem with cheaper compute, open data, OSS, and accessible education services, and (9) the growing demand for high-quality open data, fuelled by the increasing importance of data in Consumer AI development. Finally, (10) the effects of Consumer AI on labor markets were found to be less apparent. While other research typically focuses on manual occupations, we find evidence of tangible effects of AI on cognitive jobs, underlining the need for the government's existing structural unemployment strategies.

This report aims to stimulate conversation among implementers and regulators of AI-based consumer products.

Therefore, by identifying Consumer AI as a newly-emerging focus area in the field, and conducting a systematic analysis of primary research in ten interest areas, this report aims to stimulate conversation among implementers and regulators of AI-based consumer products.

INTRODUCTION

Introduction

Artificial Intelligence is one of the biggest technology interests of 2017 (Panetta, 2016; DeMers, 2016; Newman, 2017). Forbes quantifies this in financial terms, citing an expected "300% increase in investment in artificial intelligence in 2017 compared with 2016" (Press, 2016b). This trend is also supported by the increase in both public and private interest in the domain (Rao, 2017; Lunden, 2016; Simonite, 2017). Three major factors are believed to be contributing to this surge: (1) the emergence of big data, (2) more affordable access to cloud computing, and (3) easier access to improved algorithms (Kelly, 2014).

Firstly, in 2017 IBM put the recent explosion of data into perspective noting that "90% of the data in the world today has been created in the last two years alone" (IBM, 2017) . Secondly, specialised chipsets, cheap parallel computing, and scalable cloud solutions are all enabling inexpensive access to unprecedented levels of computational power (Wharton, 2017; The Economist, 2016a). Finally, easier access to algorithms is creating an ecosystem in which artificial intelligence is likely to power the next generation of consumer and business tools (Wharton, 2017; Kelly, 2014).

With an emphasis on consumer-facing technology, this report will investigate the unintended consequences of implementing artificial intelligence for personal use, within ten interest areas. It is prepared for Digital Catapult, an organisation with the mission to "drive the UK economy through the practical application of digital innovation and culture" (Digital Catapult, 2017).

This report will adopt a Methodology that is predominantly based on primary research. Semi-structured interviews were conducted with 30 experts across: business, government, and academia, regarding the topic of artificial intelligence. Furthermore, a survey of 300 consumers based in the United Kingdom (UK) was commissioned to gauge familiarity with, and sentiment towards, the topic. Our understanding of Consumer Artificial Intelligence will be informed by this primary research, along with a review of secondary sources. A working definition for Consumer Artificial Intelligence will be established to underpin our analysis and key findings.

The body of this report will present a thorough analysis of the Unintended Consequences of Consumer Artificial Intelligence in ten interest areas, as determined by primary research. Each sub-section will focus on an interest area to contextualise, discuss, and offer recommendations for policymakers, on the consumer-relevant unintended consequences. These will be further laid out in Key Findings and Recommendations, grouped in terms of direct and indirect effects on consumers. The consequences most applicable to Consumer AI will be labelled accordingly. Finally, this report will offer a Conclusion and suggest opportunities for future work, based on a holistic review of our insights.

The body of this report will present a thorough analysis of the Unintended Consequences of Consumer Artificial Intelligence in ten interest areas, as determined by primary research.

METHODOLOGY

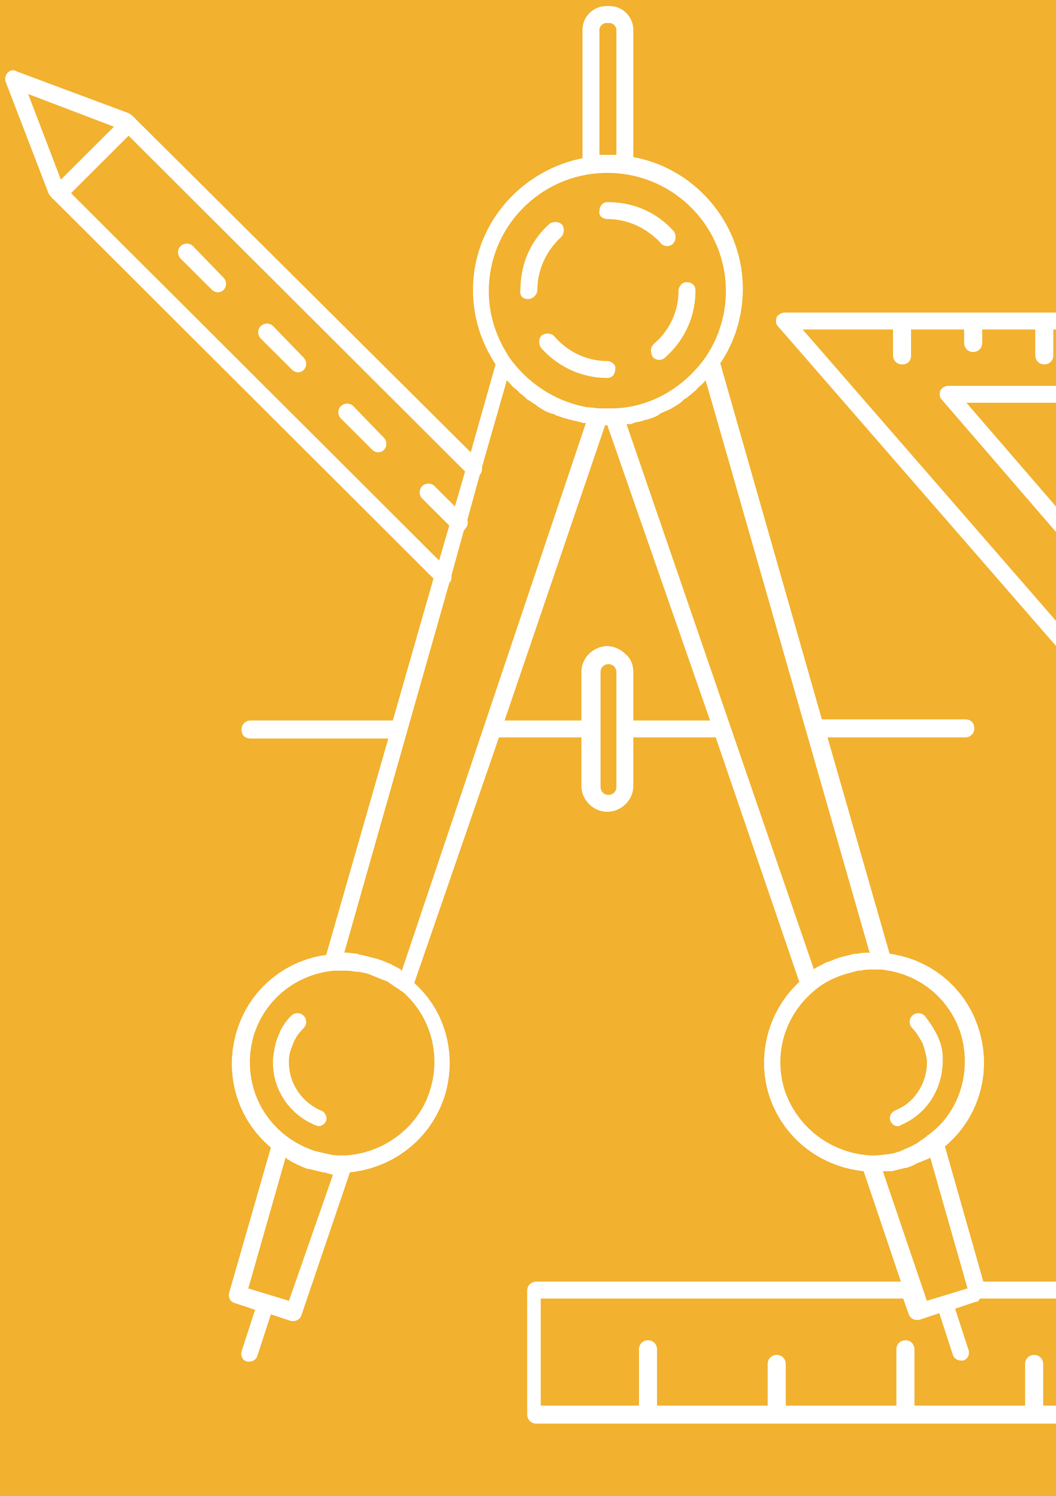

Methodology

A new wave of progress and enthusiasm in artificial intelligence has brought significant attention to economic, business, and social consequences of its developments and applications. Questions raised in these areas are being addressed from a primarily top-down, macroeconomic perspective (Agrawal et al., 2016; Smith, 2016; Amodei et al., 2016), while popular media tends to adapt a futuristic, anecdotal view on the subject. Recently published reports reveal a growing governmental interest in artificial intelligence, especially in the UK and USA (Furman, 2016; Big Innovation Centre, 2017; Executive Office of the President of the United States of America, 2017). They provide relevant, yet generic, insights into potential approaches towards artificial intelligence developments.

Despite growing interest, there is a lack of studies adopting a systematic approach towards the analysis of benefits and drawbacks beyond the intrinsic motivations of implementers. This report seeks to fill this gap by addressing the following research question: what are the unintended consequences of implementing artificial intelligence for personal use? It does so by focusing on ten interest areas that are determined by primary research. The report then provides recommendations for policymakers in each area, which also serve as reference points for business leaders and implementers.

The problem of unintended consequences of purposive actions has been at the heart of modern social and economic thought in a variety of contexts (Merton, 1936). It has been associated with the idea that certain actions always have effects that had not been originally anticipated (Norton, 2002). Following the works of Adam Smith and John Locke, the concept of unintended consequences has been associated with market and governmental failures (ibid.). Therefore, it can also provide the basis for analysis and criticism of the effects of freely operating markets and government programmes.

What are the unintended consequences of implementing artificial intelligence for personal use?

Unintended consequences are pertinent in the current era of rapidly emerging new technologies. As such, the growing popularity of consumer artificial intelligence may be expected to result in **unintended consequences**, understood as benefits and drawbacks beyond the intrinsic motivations of implementers. In this context, unintended consequences can be: predictable by applying good judgment and common sense principles ("known unknowns"), completely unexpected and for the most part unpredictable ("unknown unknowns"), or known by some but unrevealed to the public ("unknown knowns" - such as addiction triggering) (Pringle et al., 2016).

This report applies a qualitative approach to identify ten interest areas, where the growing development and popularity of consumer artificial intelligence results in unintended consequences. Key findings are informed by comprehensive primary research, based on semi-structured interviews, and a structured consumer survey.

Semi-structured interviews were conducted with 30 experts across: business, government, and academia, regarding the topic of artificial intelligence. Questions structured around two sections were answered by: 19 individuals at companies implementing artificial intelligence, 6 from governmental organisations, and 5 from academic institutions (Table 1). These organisations are based primarily in the UK, Europe, and the USA.

Companies	Governmental Organisations	Academic Institutions
• Ai Build • BioBeats • Blend • Complex • Cyberlytic • Flexciton • IBM - Watson Health • Intel • Mind the Bridge • PicsArt • Quartic • Ripjar • Skim Technologies • SoundCloud • Uber • Weave.ai • Your.MD • Zendesk + *1 Anonymous*	• European Space Agency • UK Cabinet Office • UK Department for Business, Energy & Industrial Strategy • UK Department for Culture, Media & Sport • UK Representation to the EU + *1 Anonymous*	• Leverhulme Centre for the Future of Intelligence • North Dakota University System • Oxford Internet Institute • The Alan Turing Institute • United States Naval Academy

Table 1: Individuals from the above listed companies, governmental organisations and academic institutions were interviewed for this report. An additional two organisations chose to remain anonymous. The views expressed by the individuals interviewed do not necessarily represent the views of their organisation.

The first section of the interviews aimed to build an understanding of: artificial intelligence within the given organisation and sector, the way that artificial intelligence fits into the mission or product strategy, and the broader technology policy challenges associated with the organisation's operations. The latter section of the interview aimed to provide insights into the unintended consequences of artificial intelligence applications. In order to standardise the interviews and create a robust taxonomy for the analysis, the second part of the interview was structured around four traditional market failures: monopolies, externalities, public goods, and asymmetric information (Salanié, 2000; Munday, 2000). This framework was applied as a foundation for the analysis in order to reveal higher-order unintended consequences. In this way, the report extracts and explores ten interest areas that may have beneficial or distortionary social and economic effects.

The analysis was further informed by a survey commissioned to gauge familiarity with, and sentiment towards, the topic of artificial intelligence. 300 respondents provided answers to 10 questions on their understanding, trust, and awareness of artificial intelligence in consumer products, as well as their sentiment towards the regulation and security of artificial intelligence. Survey responses were collected using *Pollfish* - a survey tool that allows for the polling of randomly selected individuals. Respondents were sampled from the population of current and potential consumers of digital products across the UK, and diversified based on sex, age, and occupation. Even though the Pollfish platform already controls for the survey bias with stratified probability sampling among the respondents, we reduced bias further by staggering survey polling into 3 iterations, each aimed at 100 respondents. These were conducted at different times of the day and week. Respondents demographics and survey results are summarised in the Appendix.

As will be discussed in the following sections this primary research was used, firstly, to provide an exhaustive definition of consumer artificial intelligence, and secondly, to inform the analysis. These are complemented by a comprehensive review of theoretical accounts, secondary sources, and case studies in the fields of economics, business, and artificial intelligence.

Finally, recommendations for policymakers in each interest area are suggested, in light of the identified unintended consequences. The recommended measures were informed by conversations with businesses, policymakers, and academics in the field, as well as secondary research of currently applied policies. In this way, this report provides informed policy recommendations in response to the unintended consequences of consumer artificial intelligence.

DEFINING CONSUMER ARTIFICIAL INTELLIGENCE

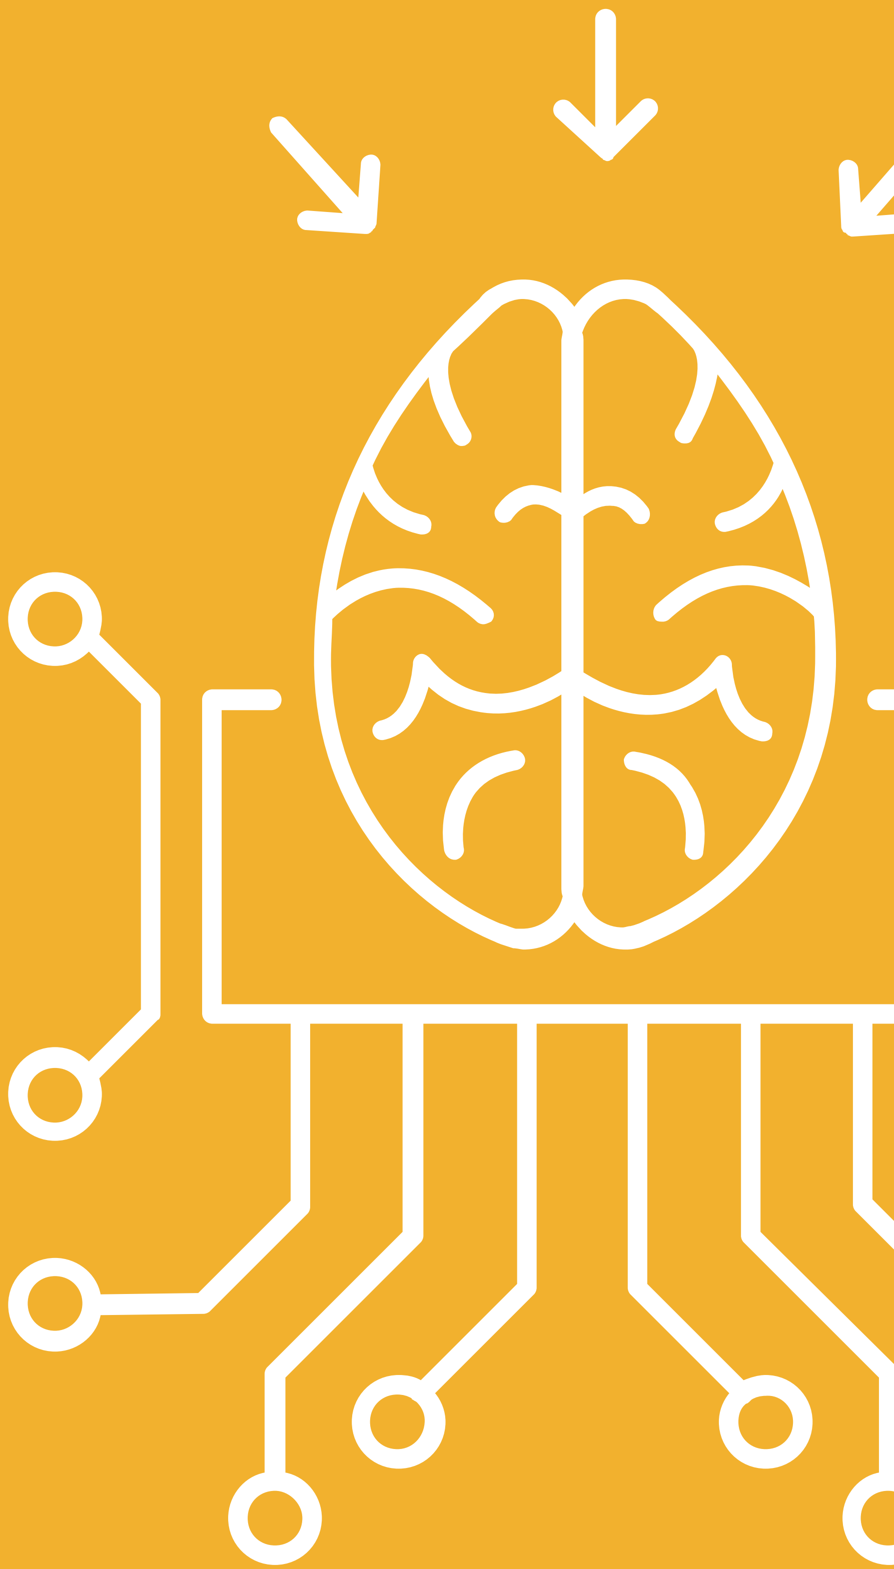

Defining Consumer Artificial Intelligence

The term **Artificial Intelligence (AI)** has a long and intricate history with roots in philosophy, mathematics, psychology, neuroscience, computer science, and linguistics. Existing definitions of AI are often confused by the discipline's various academic subfields (e.g. neural networks, machine learning, computer vision, natural language processing, etc.), as well as its portrayal in media and popular culture. It is therefore not surprising to discover that, when we asked, both public and private sector organisations claimed that the term had become overloaded with countless meanings.

Both public and private sector organisations claimed that the term had become overloaded with countless meanings.

The academic field of artificial intelligence is concerned with attempts to understand and build intelligent entities (Russell and Norvig, 2010). According to seminal work from Russell & Norvig (2010), the various definitions of AI exist on two dimensions, represented in Table 2 below. The top row is concerned with a system's reasoning process, whereas the bottom row considers a system's behaviour. Definitions to the left measure the success of a system in terms of human performance, and those to the right in terms of ideal intelligence. Ideal intelligence is referred to as rationality and reflects a system's ability to do the right thing (ibid.). These definitions lead to four possible goals for AI research, each of which has been widely pursued using different methods by different people.

	Human Performance	**Ideal Intelligence**
Reasoning Process	Thinking Humanely	Thinking Rationally
Behaviour	Acting Humanely	Acting Rationally

Table 2: Definitions of AI are organised into four categories (Russell and Norvig, 2010).

When asked about the goals of AI, organisations we spoke to for this report described systems that fell into one of two broad categories: anthropomorphic or analytical. When describing systems with anthropomorphic goals, we heard phrases such as "AI emulates human intelligence" and "AI completes human tasks". These views align with the human-centred group of definitions in Table 2, which measure success based on human performance. The second group of descriptions we heard were analytical, based on mathematics and advanced information processing. In this category, we heard phrases such as "AI is an evolution of statistical modelling" and "AI is made up of rational agents". These descriptions align with the rationalist approach to AI, which considers ideal intelligence and focuses on mathematics and engineering.

In our discussions, we also observed a loose relationship between an individual's view of AI systems and their attitude towards regulation and policymaking. Individuals who identified the goals of AI systems as anthropomorphic were more likely to talk about the technology as if it was human; suggesting an extension of existing human-centred regulations to AI systems (e.g. taxation). On the other hand, individuals who adopted a rationalist approach to AI discussed regulation at the input or point-of-action rather than of the system as a whole. In this way, they created a distinction between the underlying mathematics (known as agent function) and a system's actions.

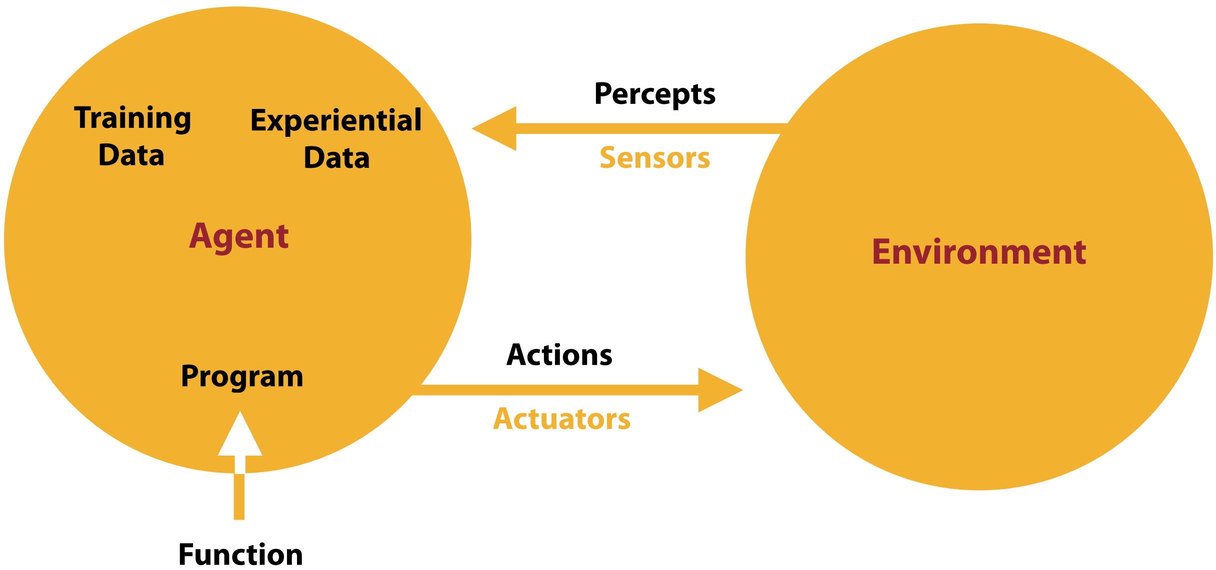

Figure 1: Agents bring together sensory input (including experiential data), with training data (where available), and an agent program (an implementation of the agent function), to act upon their environment using actuators. Adapted from Russell and Norvig (2010).

An intelligent agent is a system that can perceive its environment through sensors, and then combine this with a program and actuators to take the best possible action upon its environment.

This report will adopt a rationalist approach to AI. Intelligence will be viewed as being concerned with acting rationally "so as to achieve the best outcome or, when there is uncertainty, the best expected outcome" (Russell and Norvig, 2010, p.4). The report will focus on artificial intelligence as the study of intelligent agents; where an **intelligent agent** is a system that can perceive its environment through sensors, and then combine this with a program and actuators to take the best possible action upon its environment. All agents can come into being, and continuously improve their performance, by learning from an existing set of training data or simply through experience.

Intelligent agents are being applied to a wide range of consumer applications, making AI an increasingly significant part of everyday life. An example of adapting AI for personal use is UK-based Biobeats, who combine health data from wearables with evidence-based machine learning, to help consumers manage their health and productivity. Skim.it is a Google Chrome extension for students and teachers that uses AI to extract key information from articles on major news websites. Additionally, Jukedeck, with origins at the University of Cambridge, uses AI to compose and adapt professional-quality music, giving consumers audio files that are personalised to their needs. Consumer AI also exists beyond the realm of software agents, playing an increasingly important role in consumer electronics. As an example, the world's most successful robotic vacuum cleaner, the Roomba, adopts AI techniques to cover surface areas more efficiently (Knight, 2015).

The emergence of personal assistants such as Apple's Siri, Microsoft's Cortana, Amazon's Alexa, and Google Assistant, are all examples of AI deployed for personal use. What is more, SoundCloud relies on machine learning to determine what songs to recommend next, while Netflix, Facebook, and Snapchat use subfields of AI such as computer vision, to improve their core products.

Consumer AI is defined as the commercial development of intelligent agents for personal use.

Informed by the discussion on artificial intelligence, we have coined the term Consumer Artificial Intelligence to describe this empirical trend. **Consumer AI** is defined as the commercial development of intelligent agents for personal use.

In addition to these empirical examples of Consumer AI, our survey shows that the three most used technology categories among consumers are ones that make extensive use of AI: social networking (e.g. spam filtering), online shopping (e.g. product recommendations), and entertainment (e.g. computer vision applications). This ongoing development and adoption of Consumer AI is likely to lead to unintended consequences. This report explores unintended consequences in ten of the most commonly discussed interest areas, which emerged from our conversations with businesses, policymakers, and academics.

UNINTENDED CONSEQUENCES OF CONSUMER ARTIFICIAL INTELLIGENCE

Unintended Consequences of Consumer AI

The primary research conducted for this report sought to explore and understand potential unintended consequences resulting from the implementation of Consumer AI. With this approach, ten interest areas were identified, some of which had special significance in their relationship to Consumer AI; others were applicable to the AI field more broadly. This section explores the ten interest areas, by examining the consequences in each and thus making appropriate policy recommendations.

1. Consumer Knowledge

The level of understanding of technology among consumers of AI-specific applications was one of the key interest areas of this report. It appears that the majority of consumer survey respondents overstated how knowledgeable they were about new technologies, despite 9 in 10 identifying themselves as at least somewhat tech-savvy.

4 in 10 stated that technologies they use take advantage of AI, however only a quarter of those (1 in 10) said that they could name a specific application. This is despite the fact that social networks, shopping, and entertainment were among the most popular application categories used by consumers; all of which increasingly take advantage of AI. These findings align with the academic understanding of overconfidence as "the overestimation of one's actual ability, performance, level of control, or chance of success" (Moore and Healy, 2008).

This, along with the findings derived from conversations with companies, informs two considerations. Firstly, consumers consider AI in terms of a "black box" and lack knowledge about its inner-workings. Secondly, and more importantly, consumers appear not to know what they do not know about AI. These AI-specific "unknown unknowns" can be considered a risk factor associated with Consumer AI, which the subjective decision-maker does not imagine and therefore does not even consider (Feduzi and Runde, 2014).

The discrepancy between people's perceived and actual understanding of technology is growing with the deployment of Consumer AI.

In summary, the organisations we spoke to for this report have revealed that the discrepancy between people's perceived and actual understanding of technology is growing with the deployment of Consumer AI. They also expressed a clear desire for actions which could reduce this discrepancy. These sentiments were further reflected in our consumer survey, where 9 in 10 respondents revealed their desire for more information about how AI technologies work. Thus, there is an opportunity for both the private and public sector to provide more information on the workings of AI-reliant products and services. Our research revealed that many businesses have already recognised this need, and are addressing it by introducing transparency that enhances the user experience. This will be discussed further in the section on the Principal-Agent relationship.

> There is an opportunity for both the private and public sector to provide more information on the workings of AI-reliant products and services.

2. Principal-Agent Relationship

According to traditional agency theory, the so-called agency problem occurs when cooperating parties (the principal and the agent) have different goals or interests, and the agent's behaviour cannot be verified appropriately by the principal (Eisenhardt, 1989). Agency theory has been considered by scholars in a variety of academic areas, mostly in economics and corporate finance (Laffont and Martimort, 2009; Nikkinen and Sahlström, 2004; Li, 2011). What is more, the agency problem, particularly in terms of a general theory of the principal-agent relationship, has been applied to a range of interactions, for example: employer-employee, lawyer-client, or buyer-supplier (Harris and Raviv, 1978).

Interestingly, our primary research revealed that the growing popularity of AI applications may see increasing prevalence of this traditional agency problem in the Consumer AI domain. In conversations with implementers of AI and key influencers in the field, we found that this agency problem results primarily from a lack of transparency between the inner motivations of an AI agent and the principal (i.e. the consumer). Examples of such agency conflict may appear in recommendation engines such as Amazon (e.g. "Your Recommendations") and Spotify (e.g. "Discover Weekly"), as well as in functionalities of popular intelligent assistants such as Amazon's Alexa and Apple's Siri.

It has been suggested to us that AI implementers may have an incentive to develop agents whose interests do not align with those of the principal, but instead primarily serve their business interests. Therefore, questions of agency transparency, traditionally understood as adverse selection, can arise. The agent is assumed to possess certain abilities, however the principal cannot completely verify these (Eisenhardt, 1989). This can be illustrated with the following example of a fictitious recommendation engine.

We shall consider a consumer (the principal) who uses a recommendation engine based on an AI agent, on a platform that recommends products for purchase. It is in the consumer's interest to have the product recommended, based on input data related only to the principal, such as purchase history, personal demographics, and other consumer-specific characteristics. However, the developer of the recommendation engine may be incentivised to recommend a product, which instead of serving the consumer's interest, furthers the platform's business interests. An example of this would be recommending a product that is in favour of the business's bottom line. Therefore, in the Consumer AI realm, the agency problem arises when the balance between the platform's business interest and the consumer's interest is skewed.

A proposed solution to the Consumer AI agency problem may be to increase transparency about an agent's inner motivations, creating a balance between business and consumer interests. Businesses may choose to potentially compromise their competitive advantage by revealing their approach, or a part thereof (e.g. confidence levels), on how recommendations were derived. This would be justified as an attempt to address the stated agency problem. Practical examples of this level of transparency have been demonstrated by Amazon's move to open source its recommendation engine DSSTNE (Novet, 2016) and Netflix's "Watch Next" feature, which discloses recommendation confidence levels. Our primary research, however, reveals that despite steps such as these, low levels of understanding regarding recommendations and confidence levels persist among users. Therefore, questions remain as to the effectiveness of this approach in balancing consumer and business motivations.

The traditional principal-agent problem sees a growing prevalence in applications of Consumer AI.

In summary, while the traditional principal-agent problem sees growing prevalence in applications of Consumer AI, increased transparency alone does not necessarily correct for the misalignment between commercial and consumer interests. We recommend that governments avoid forcing transparency requirements as this has been shown to be ineffective, and we suspect that it may lead to an unnecessary loss of competitive advantage for businesses concerned. At the same time, the distinct nature of some AI agents makes their actions difficult to predict or explain, even by their implementers; this agent scrutiny problem is discussed in a later section.

We recommend that governments avoid forcing transparency requirements as this has been shown to be ineffective, and we suspect that it may lead to an unnecessary loss of competitive advantage for the businesses concerned.

3. Agent Scrutiny

Our primary research revealed that certain AI techniques make it increasingly difficult to confidently explain and predict the behaviour of some AI agents. However, it is not only the consumer that may experience problems with verifying an agent's behaviour, as in the principal-agent problem already discussed. At the current level of AI complexity, system designers themselves cannot foresee consequences that may results from some of their applications (The Economist, 2015). This phenomenon can be referred to as the agent scrutiny problem, where even the engineers who build these applications cannot fully explain their behaviour (Knight, 2017).

Increased use of certain AI techniques leads to the agent scrutiny problem, and thus requires a distinct regulatory approach that does not impede AI research.

At the time of writing this report, an example of an AI technique that can result in the agent scrutiny problem is Deep Neural Networks (DNNs) (Lee et al., 2015). On this point, an increased use of such AI techniques leads to the agent scrutiny problem, and thus requires a distinct regulatory approach that does not impede AI research.

We find that the agent scrutiny problem in Consumer AI has raised concerns among academics and policymakers. The difficulty in scrutinising the agent has found its way into public debate, recently and notably by the German Chancellor Angela Merkel.

In response to the spread of fake news and hate speech on online platforms, in October 2016, Merkel's administration publicly requested that major internet platforms make algorithms they use open to scrutiny. She stated, "I'm of the opinion that algorithms must be made more transparent, so that one can inform oneself as an interested citizen about questions like 'what influences my behaviour on the internet and that of others? [...] Algorithms, when they are not transparent, can lead to a distortion of our perception, they can shrink our expanse of information" (Connolly, 2016).

In the UK, similar sentiment was shared by the Labour Shadow Minister for the Digital Economy, who in December 2016 called for the algorithms used by technology firms to be made transparent and subject to regulation. She wanted to see greater scrutiny of the "mathematical formulas" that control everything from the tailored news served to Facebook members, to the speed at which workers are required to move around an Amazon warehouse (Garside, 2016). She also underlined that: "Algorithms are part of our world, so they are subject to regulation", while recognising that "because [algorithms] are not transparent, it's difficult to regulate them effectively" (ibid.).

These statements represent a desire among policymakers and citizens to better understand how AI algorithms work. Such transparency requires a level of understanding that is difficult to attain by even the developers themselves (Knight, 2017). It demonstrates that there remains a lack of understanding among policymakers as to how certain AI applications work. What is more, regulation at the agent level may lead to inefficient restrictions on agent function research or agent program development.

Our primary research revealed that applications of AI are commonly viewed as a source of competitive advantage for companies. The majority of those we spoke to for this report agreed that regulating at the agent may hinder the progress of research and development in the field. The spread of fake news and hate speech, which served as catalysts for these calls, are certainly a drawback of Consumer AI. However, we believe that the proposed form of government intervention, which demands more "transparency" of the agent is a flawed approach because of the aforementioned agent scrutiny problem. An alternative intervention would focus on restricting the publication of objectionable content (e.g. hate speech, as the input to a recommendation agent) and the surfacing of said content (e.g. recommended articles as the action of the agent). Thus in this illustration, the agent function research and program development are unimpeded.

It is important to note that there are countless useful applications of AI (Yao, 2017). Some such applications that also suffer from the agent scrutiny problem, range from language translation software to image recognition; all of which have increased need for Consumer AI. We believe that the research and development of AI techniques should remain uninterrupted such that implementers can continue to provide consumers with useful products and services.

In addressing the agent scrutiny problem, we recommend that if governments decide to intervene on matters relating to Consumer AI, such as in the case of hate speech, the focus should be on an agent's inputs and actions rather than the function or program. This novel concept of agent decoupling enables regulation that does not impede research and development into Consumer AI.

We recommend that if governments decide to intervene on matters relating to Consumer AI, such as in the case of hate speech, the focus should be on an agent's inputs and actions rather than the function or program.

Case Study – Deep Neural Networks

Deep Neural Networks (DNNs) are seen as "the hottest" topic in speech recognition, computer vision, and natural language processing (Marc'aurelio and Lecun, 2013). A high profile use of DNNs was that of AlphaGo, created by DeepMind. By playing "a handful of highly inventive winning moves, one of which was so surprising it overturned hundreds of years of received wisdom", AlphaGo beat an 18-time world champion Go player (Hassabis, 2017). This example shows the potential of DNNs in pushing the boundaries of AI, however it also highlights the agent scrutiny problem: "can we explain how these agents work?"

DNNs are built, to some degree, to mimic the neurons that power human intelligence. The architecture for such systems is based on three layers of nodes (neurons): input, hidden, and output layers (see Figure 2). Data enters through the input layer, then travels, based on mathematical weights, through one or more hidden layers, before ending at the output layer (Gershenson, 2003). A characteristic of DNNs is that implementers, generally, can only decipher the function's decisions at the input and output layer. The processes and decisions made in the hidden layers are often unsupervised and complex, due to the changing value of weighting that constantly adapts for better results (Russell and Norvig, 2010). Thus, the complexity of scrutiny for such systems can be problematic.

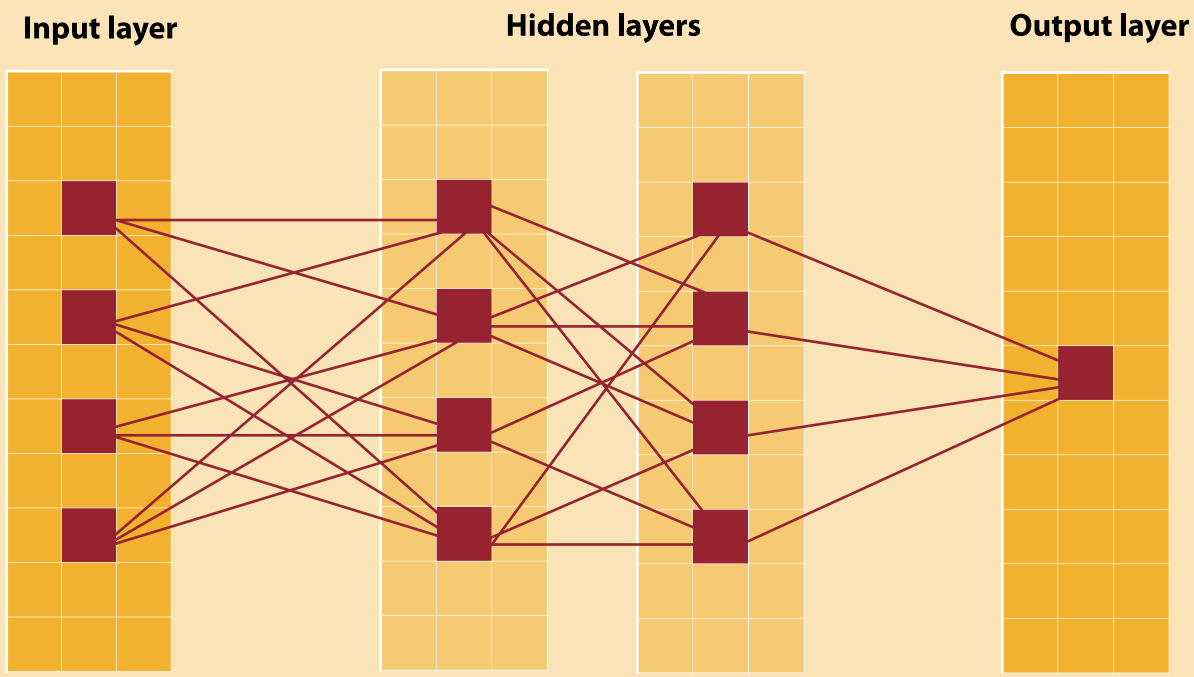

Figure 2. In Deep Neural Networks, data travels from the Input layer, through one or more Hidden layers, before ending at the Output layer.

4. Discrimination

A concern mentioned by many of the organisations consulted was that prejudice in the development and training of intelligent agents can lead to codified discrimination in Consumer AI. While AI can act consistently and objectively, recent cases reported in the media have shown how unjust behaviour may manifest. For example, Google's facial recognition software was recently found to accidentally categorise a photo of two African-American individuals as gorillas. Other incidents include Flickr tagging photos of traditional Native-American dancers as "costume" and Nikon asking "Did someone Blink?" after a photograph was taken of an Asian-American person's face (Zhang, 2015). The aforementioned cases are warnings of ways Consumer AI may marginalise certain groups in society. This section offers a non-exhaustive list of measures companies can take to reduce risks of discrimination within Consumer AI products.

Prejudice in the development and training of intelligent agents can lead to codified discrimination in Consumer AI.

Data treatment is one method to lessen codified bias in Consumer AI. An example of this is the World White Web, a project started by Johanna Burai when she found that an online image search of the word "hand" produced results with only white hands (World White Web, 2016). These results were a function of the frequency and popularity of "hand" images across the web, which were not representative of true global population demographics. The algorithm attempted to serve her the most relevant content, but in doing so exhibited discriminatory behaviour. Her project seeks to encourage users to share images of multi-ethnic hands across the web to retrain the agent. Rather than retroactive action, companies may refer to this as a lesson for intentionally using diverse training data, from the onset of development, to prevent prejudice.

Beyond data centric solutions, companies have a range of other measures that they can put in place to avoid codifying discrimination. The first is targeting the homogeneity of the workforce itself. Many large technology firms lack widespread diversity (Stacy Jones, 2015). This lack of diversity can cause development teams to unintentionally bias the algorithms they create (Garcia, 2017). Code for America has a project to address this by evaluating and highlighting instances of discrimination in technology job posts (Joblint, n.d.). An article in Nature Scientific Journal revealed another approach which includes user input into the learning model (Maher, 2016). Such feedback, when applied directly to discrimination, can help train intelligent agents to see things that they may not otherwise. Finally, research from Google offers yet another solution to "attacking discrimination with smarter machine learning" by focusing on "threshold classifiers" (Martin et al., 2016; Hardt et al., 2016). A complete analysis of all proposed solutions is outside the scope of this report, yet it is important to note some of the many methods companies may utilise.

We recommend a review of available options by both government and businesses, to reduce risks of codifying discrimination in implementations of Consumer AI.

Some have even suggested governmental "certification for companies actively and thoughtfully working to reduce algorithmic discrimination" (Garcia, 2017). Academic work has also suggested governmental involvement to address regulatory voids in existing legislation to account for a lack of "transparency and accountability of automated decision-making" (Wachter et al., 2016). We recommend a review of available options by both government and businesses, to reduce risks of codifying discrimination in implementations of Consumer AI.

Case Study – MeVitae is leveraging AI to tackle workplace discrimination

MeVitae is an AI-based technology solution that helps identify unconscious bias. It uses a range of sensorial data captured while employers review Curriculum Vitae, to algorithmically enable "bias correction". In doing so, MeVitae hopes to break down the barriers "hindering diversity in the workplace".

Organisations can combine tools such as MeVitae with strategies that minimise risks of data or algorithmic bias, to prevent the codification of discrimination within their Consumer AI offerings (Doraiswamy, 2017).

5. Cyber Security

Few discussions on technology exist void of cyber security considerations; AI is no exception. Our research revealed concerns about the heightened impact of data breaches, compromised agents, and offensive capabilities. Consequently, we recommend that every organisation consider Consumer AI as a core component of their cyber security strategy.

> Our research revealed concerns about the heightened impact of data breaches, compromised agents, and offensive capabilities.

Data protection was by far the most widespread and frequent. In order to build useful AI models, algorithms typically learn from massive training data. Furthermore, when designing AI for consumer use the nature of this data is inherently personal. Capturing, analysing, and then acting on this data is at the centre of Consumer AI applications. As such, a large number of organisations identified that data protection was a key element to their security strategy.

Companies noted that some data protection regulation already exists to cover consumer technologies. They cited existing consumer privacy protections such as the UK Data Protection Act (1998) and the European Union's General Data Protection Regulation (2016), underpin their security measures.

Yet, data breaches remain prevalent (Ponemon Institute, 2016). An economic problem exists as companies consider how to invest in protection against these developing threats.

Digital security follows a network-based benefit model. For example, as more people choose to use antivirus software, the spread of vulnerabilities may decrease across the network and in turn raise the overall digital ecosystem's defensive posture (Kunreuther and Heal, 2003; Camp and Wolfram, 2000). This distributive security model creates positive externalities resulting in "sub-optimal investment" (Kunreuther and Heal, 2003; Camp and Wolfram, 2000) and it "changes the maximum a firm should, from a social welfare perspective, invest in cyber security activities" (Gordon et al., 2015). Our conversations reflected this research, as many organisations claimed that cyber security defence systems had high costs.

> We recommend that every organisation consider Consumer AI as a core component of their cyber security strategy.

We observed that this sentiment was proportional to the size of the company in question. When compared to larger firms, start-ups claimed to have less resources available for cyber security. Some research suggests that economic intervention, proportional to a firm's size, may encourage more optimal investment in cyber security (Gong et al., 2009).

Potential brand damage, resulting from data breaches, is another key factor in security investment calculations. A Ponemon Institute (2012) study noted that 41% of companies who experienced a data breach also experienced loss of customer loyalty. Our survey validates this consumer worry, as an overwhelming majority of respondents showed concern for the security of their data. Another study calculated that loss of customer data was the most costly type of data breach, when compared to loss of employee data or intellectual property (Ponemon Institute, 2011). Combining increased brand damage costs, with heightened risks of holding copious personal data, highlights the unique impact of data breaches in Consumer AI.

The second trend that emerged in our research was that of compromised agents. As highlighted by the agent scrutiny problem, some Consumer AI behaviour cannot be explained, even by implementers. Organisations we spoke to expressed concern that implementing such AI techniques may prevent them from detecting malicious behaviour until the whole model has become corrupted. The Microsoft Tay chatbot illustrates a manifestation of this corruption. Tay was built to simulate an innocent 16-years-old girl on Twitter, interacting in a colloquial, relevant manner. Within hours of deployment, the Tay bot became a vehicle of harassment and hate speech, as a result of unfiltered data corrupting the agent. Microsoft quickly took Tay offline for "upgrades" and deleted the most controversial tweets (Price, 2016).

Organisations cited such corruptibility as a key concern for Consumer AI reliant on "black box" algorithms. Some highlighted the importance of data treatment as an aid to this problem, however this solution comes with a large capital burden. Another suggestion was the fundamental decoupling of an agent's action from its program. Applying this to Tay, could have taken the form of a human in the loop. In such a scenario, the human would sit between the agent and its environment. The human would be responsible for approving data coming into the agent and any actions taken as a result. Admittedly, such oversight undermines key motivations of automation and efficiency achieved through AI in the first place. Companies should consider the potential ramifications of their applications of

Consumer AI and adjust the independence of agents accordingly. Solutions for this were noted in the section on agent scrutiny.

The third trend that emerged from our research was the adaptation of AI to enhance the offensive capabilities of malicious actors. The development of Consumer AI in the field of cyber security has typically involved securing personal systems. For example, Symantec offers machine learning enhanced endpoint protection for personal devices (Symantec, n.d.). However, misappropriation of these defensive systems is evident in cases such as the most recent DARPA Cyber Grand Challenge (DARPA, 2016). Teams competed to create the most effective AI enabled cyber security capabilities to detect vulnerabilities in their own systems, while exploiting those on other team's (Coldewey, 2016). Former President of the United States, Barack Obama, shared this concern, surmising that AI enabled hacking may lead to catastrophes as dire as the uncontrolled launch of nuclear rockets (Greenberg, A, 2016). From this extreme case to a simple increase in more common cybercrime, regulation must be updated and bolstered to account for new offensive capabilities.

6. Market Structure

Our research found that the market dynamics of Consumer AI may lead to market dominance, however there is no consensus regarding anti-competitive behaviour. As defined previously, intelligent agents can learn from two sources, existing sets of training and experiential data. Individuals we spoke to emphasised this point, noting that access to large amounts of data was "key" to developing Consumer AI. Publications such as The Economist echo the importance of data in AI algorithm development (The Economist, 2017a). In analysing the emergence of market leaders in the space, it is therefore necessary that we consider proprietary data as a core topic.

> The market dynamics of Consumer AI may lead to market dominance, however there is no consensus regarding anti-competitive behaviour.

Proprietary training data is particularly important to companies developing Consumer AI due to the personal nature of the data being collected. Nearly any product consumers engage with, from smart TVs to smartphones, generates personal data regarding their demographics, preferences, and behaviours (The Economist, 2017b). This is particularly useful in Consumer AI products because they are created to cater to personal needs. The personal nature of the data makes it a valuable asset to companies, but it also involves significant risk if the data is in any way compromised (O2, 2015). To extract valuable insights into consumer behaviour and preferences, the topic of proprietary data goes beyond the possession of data, as was stressed by those we spoke to. To make the data valuable, organisations must also structure and process the data in ways that require the finances, infrastructure, computational power and human capital that are typically more available to large companies (The Economist, 2017b). Large companies that have these capabilities can apply their resources strategically to gain a significant competitive advantage (Weinberger, 2016).

The market dynamics of companies implementing Consumer AI provide further insight as to how companies solidify their market leadership. Software products, and "new economy" information technology industries, are by their nature and use-cases subject to network effects (Farrell and Klemperer, 2006). "Network effects" is the idea that products increase their inherent value as their user-base expands (Gallaugher and Wang, 2002). An example of network effects that are expedited by Consumer AI can be seen in Facebook. As more users adopt the service, the utility of the service increases for each individual user. Thus, an important portion of the value of Facebook is in its network, since the data generated from each additional user allows the product to improve, and as its functionality improves, so does the user experience. Hence, even a company with the technological and human capital necessary to compete with Facebook, would be faced with the challenges posed by these winner-takes-all (or winner-takes-most) network effects (Kemerer et al., 2013) . Experts consulted for this report acknowledged that certain market leaders had already emerged as a result of these market dynamics. Consequently, they raised concerns that the increased deployment of Consumer AI, may lead to environments that enable monopolistic behaviour.

In our primary research, we identified two opinions about the potential of market dominance. The first opinion, held primarily by larger companies we spoke to, was that the novel nature of digital market dynamics has not yet yielded any true monopolies in the traditional economic sense. Although proprietary data can give companies significant advantage, the opinion followed that the nature of consumer generated data makes it difficult to consolidate into the hands of one single company. What is more, regularly mentioned was the point that data availability and accessibility do not, on their own, lead to a competitive advantage. Rather, and more importantly, it is how the data is used. When it comes to data usage, some individuals we spoke to believe that open source software can provide the support newcomers require to build competitive Consumer AI applications. Finally, some individuals suggested that existing laws pertaining to competition, data protection, and consumer rights, already provide adequate regulatory frameworks.

The second opinion was more precautionary, expressed primarily by smaller companies and startups. A belief emerged that while traditional monopolies may not yet exist, there is clear potential for them to form in the near future and "modern monopolies" may emerge. Startups shared experiences of consciously pivoting their business away from areas where they suspected imminent competition with large companies. In the case of data collection and application, companies expressed a sense of being "deterred" by the fear of "never being able to catch up" with market leaders, calling this an "anti-competitive environment." Companies also emphasised the importance of public opinion on the matter, and in our survey to gauge sentiment, almost half of respondents acknowledged the importance of Consumer AI for corporate competitive advantage. These findings, as expressed herein, align with recently published work discussing the nature of a data-driven digital economy and the risks of data concentration (The Economist, 2017b).

Given the importance of personal data in Consumer AI, as illustrated above, we believe that the issue of proprietary data consolidation is understudied; especially when viewed through a traditional antitrust lens. Certain individuals expressed desire to address the consolidation of data, as we do with that of tangible commodities such as oil. This may be applicable in some merger and acquisition cases, however, the liabilities associated with personal consumer data make it difficult to view this activity as an asset consolidation in the traditional sense (O2, 2015). This requires critical rethinking of current antitrust policy and some have already proposed regulatory approaches to data concentration (The Economist, 2017b). As non-regulatory approaches, Solid and CitizenMe seek to give internet users control or ownership over their own data, reducing the potential for monopolistic gains (The Economist, 2017a; Weinberger, 2016).

In summary, it is difficult to reach an exhaustive conclusion regarding the topic of market dominance, or the potential for anti-competitive behaviour in the age of

Consumer AI. What we have observed is that despite the network effects discussed and the role of Consumer AI in propagating them, currently there is no consensus on the issue of market dominance in Consumer AI among businesses, policymakers, or academics. We recognise that this is an area of research that continues to demand the attention of experts in technology, law, economics, ethics, and beyond. Hence, our recommendation for policymakers is to continue studying the issue, and adopt a novel competition lens in the spirit of ensuring that the development of Consumer AI continues to serve consumer interests first - intervening only in cases of anti-competitive behaviour.

Our recommendation for policymakers is to continue studying the issue, and adopt a novel competition lens in the spirit of ensuring that the development of Consumer AI continues to serve consumer interests first - intervening only in cases of anti-competitive behaviour.

7. AI Research

Our research found that there is a shift of academics to the private sector and hence the privatisation of future AI research. This is related to two trends. The first, and most important, was the growing trend of AI experts and researchers leaving academia for jobs in the private sector. The second, was the private acquisition of AI intellectual capital within open source communities.

There is a shift of academics to the private sector and hence the privatisation of future AI research.

The first trend marks a potential shift away from how AI has historically evolved as a field, driven primarily by academia. Drawing on previous work by Alan Turing, the term artificial intelligence was coined by a Dartmouth academic, John McCarthy in 1956 (Press, 2016a). AI research in academia has fostered an environment in which progress in the field has been openly shared as public knowledge. This open environment however is changing, as there is a body of evidence showing that researchers are moving from academia to the private sector at an uncommon rate. At least three examples of this that can be drawn from the Consumer AI domain. In 2013 Facebook recruited Yann Lecun from New York University to head up their AI lab

(Kagel, 2013). In 2015, Uber hired 40 AI and robotics staff and researchers from Carnegie Mellon University (The Economist, 2016c). In 2016, Google was able to recruit the leader of Stanford University's AI lab, Fei-Fei Li (Vanian, 2016).

Academics exiting academia to work for big technology firms may be associated with attractive compensation packages and bonuses (Royal Society, 2017). Another potential cause discovered in our research and corroborated by a recent article in The Economist, is the ability for researchers to focus on their work without distractions of securing funding (The Economist, 2016c). Organisations we spoke to also noted the ability of private companies to provide researchers with proprietary datasets, larger engineering teams, and advanced facilities.

There are risks associated with the movement of academics, and we will focus on two. Firstly, we believe that the "brain drain" of academics to industry may hinder the progress of academic research in the AI area (Waters, 2016). Students of AI may suffer as prominent academics leave the field, taking with them the supervision and guidance necessary to foster new talent (Royal Society, 2017). Further, graduate students themselves may leave the field of academia prematurely in pursuit of lucrative jobs at technology firms (The Economist, 2016c). Secondly, with the movement of academics to industry, new research may not be made available to the public. This can lead to companies developing significant competitive advantage, due to the coupling of massive proprietary datasets with the enhanced ability to develop improved proprietary algorithms. Companies can use this to achieve market leadership in the realm of AI.

We recommend enhanced investment by government to ensure academia remains an attractive environment for researchers to advance the field of AI.

The second trend is the acquisition of open source, research-heavy, AI startups by private companies. Kaggle, a "platform that hosts data science and machine learning competitions" was acquired by Google in 2017 (Lardinois et al., 2017) and in 2015, Facebook acquired Wit.ai thereby acquiring a community of over 6000 developers (Constine, 2015). In both cases, companies expressed that the communities, as well as the algorithms developed, would remain free and open. However obvious risks remain. Should the acquirer change their open source strategy, or choose to keep the most significant discoveries proprietary, the wider AI development community would clearly stand to suffer.

This would not be unprecedented, as in 2013 after the acquisition of PrimeSense, Apple proceeded to privatise the associated open source computer-vision technology: OpenNI, impacting 100,000 active developers (Armstrong, 2014).

The shift of academics into the private-sector and acquisition of open source communities are not by themselves cause for concern. What may be however, is the potential abuse of the market leadership that can be gained from this acquisitive behaviour. To address these concerns, governments may choose to increase funding for AI research, thereby combating the drain of academics into industry. Recently, the UK government committed £17.3 million in funding for AI research at universities in the UK (Murgia, 2017). This is a positive step in the right direction, however we recommend enhanced investment by government to ensure academia remains an attractive environment for researchers to advance the field of AI.

8. AI Developer Ecosystem

While gathering our primary research, Open Source Software (OSS) was a topic raised by most companies, as it is an essential component of their technology stack. OSS allows for the (typically free) licensing of technology to distribute, modify, or commercialise software in a manner that enables open collaboration (Perens and Others, 1999). Implementers we spoke to claimed to have a heavy reliance on open source AI development tools, frameworks, and libraries. This was especially evident in conversations with early stage companies. Not only do these businesses rely on open source AI projects, but many actively contribute back to them. This interest and dedication in the open source AI community, makes it vibrant and beneficial for all involved.

Several of what were previously proprietary AI projects, have now been released as open source. Google, Facebook, and Amazon have all released machine-learning platforms, these are TensorFlow, Torch, and DSSTNE respectively (Finley, 2016a, 2016b). While the endorsement of OSS may, on the surface, seem irrational from a traditional economic perspective, it is important to note the positive externalities mentioned: attracting and retaining talent, improved developer community image, and the

Collectively: cheaper compute, open data, OSS, and accessible education services, are democratising the AI developer ecosystem.

ability to own and influence industry standards.

We heard from businesses that the adoption of OSS by large companies is having a meaningful effect on the development of not only Consumer AI applications, but the AI developer ecosystem overall. For instance, TensorFlow is an OSS library for machine intelligence, which provides advanced tools to developers who may have never had access to this state-of-the-art technology otherwise. A recent application of TensorFlow is that of a Japanese farmer, who with no previous AI development experience, was able to build a working application within only a few months. Using a personal computer and amateur photos as training data, he achieved a 70% recognition efficiency in produce categorisation for 21 different types of cucumbers (Sato, 2016).

The only warranted government action is positive incentives that increase engagement in the AI developer ecosystem.

The AI ecosystem more generally is becoming increasingly accessible. A developer can now combine open data, with open source AI libraries, running on more affordable cloud services, to build novel AI applications (Wharton, 2017). What is more, several online services are beginning to offer free introductory classes and online resources, to help developers lift their skills and thus diffuse AI development (Mannes, 2017).

This greater accessibility to AI has materialised in the UK economy, with the number of new AI startups doubling in the three-year period between 2014 and 2016 inclusive, as compared to the three years prior (Kelnar, 2016). Collectively: cheaper compute, open data, OSS, and accessible education services, are democratising the AI developer ecosystem. This democratisation is accelerated by the Consumer AI community, therefore the only warranted government action is positive incentives that increase engagement in the AI developer ecosystem.

9. Open Data

As was unanimously expressed in our discussions, the implementation of Consumer AI has resulted in an increased demand for quality data, as an integral part of developing and improving AI applications. Quality training data is often challenging to gather and requires resources to structure. An often suggested solution was increased availability, usability, and awareness of public datasets. The value of open data is emphasised by publications from Deloitte Analytics (2012) and the Open Data Institute (2017). Deloitte specifically stated that "every industry will benefit from using [open data] to improve the quality, completeness, and utility of their own data" (Deloitte Analytics, 2012, p.5).

The implementation of Consumer AI has resulted in an increased demand for quality data, as an integral part of developing and improving AI applications.

Our primary research yielded a divide with regards to the importance of open data. Some organisations noted that such data was critical to their foundation, while others used little to none. Not surprisingly, smaller firms found open data to be "more critical" in their priming of Consumer AI systems, whereas larger companies reported a strong reliance on proprietary data that was more structured and tuned to their application. In summary of this point, it is generally accepted that open data is critical for the development of Consumer AI, and its value was especially emphasised by startups.

In the last few years, the UK government has made access to public data centric to its transparency agenda (HM Government, 2016). Currently, all departments are required to release their data, and datasets on: health, education, transport, and justice are already available (UK Cabinet Office, 2012). Additionally, each government department must include open data schemes in their strategies to further increase data transparency and access (ibid.). The massive datasets generated by the UK's initiative have created a large repository available to implementers as discussed in the case study. The success of data.gov.uk and other similar national databases, like data.gov and data.gouv.fr, make a convincing case for continued public investment in public data.

Creating and sharing data publicly is not unique to government, businesses of all sizes are also making contributions to open data repositories (Open Data

Institute, 2017). Inherently, the importance of data in Consumer AI development is increasing demand for high-quality open data, which can be further supported by government. Consequently, promoting an increase in the availability, usability, and awareness of public data must be a key consideration for policymakers going forward.

We offer no recommendation on the matter of data-sharing, since it is outside the scope of this report. Our hope instead, is that we have made clear the importance of public and open data, and that this will spur future research and favourable policy.

Promoting an increase in the availability, usability, and awareness of public data must be a key consideration for policymakers going forward.

Case Study – We expect Consumer AI to create significant benefits for society and the UK is well-positioned to capitalise on these technologies

Data.gov.uk, launched in 2010, is a growing example of public data as a source for large open data projects (data.gov.uk, n.d.). Today, over 40,000 unique datasets are available online to be shared and selected for implementation. This initiative and other public UK sources (e.g. London Datastore) have led to the UK being ranked first in the world for open data (Open Data Barometer, 2015). As a public good, public data can help build a robust digital future (Shadbolt, 2015; Summers, 2016). Despite criticism that the data may not be easy to access, easy to find, or structured to suit specific needs, the UK stands to benefit greatly from ongoing investment in the platform.

10. Labour Market

All organisations we spoke to acknowledged the impact of AI in the labour market. The discussion generally focused on "productivity gains" or "force-multiplier effects" in manufacturing and manual jobs. These observations mirror long-standing sentiments in academia and the public sector. For example, a recent Brookings Institute study shows that routine manual occupations are at the highest risk of automation on account of technological advancements (West, 2015). This trend is further exacerbated by the introduction of AI. A global survey by Weber Shandwick in partnership with KRC Research, revealed that 82% of markets surveyed showed concern that AI will amplify job losses (Weber Shandwick, 2016).

Interestingly, our primary research showed that the specific consequence of Consumer AI in the labour market is less apparent as it affects cognitive jobs more than the typically considered manual occupations; this creates risks of neglect or flawed response. Our finding is supported by the opinion that while "attention around automation focuses on how factory robots and self-driving cars may fundamentally change our workforce [...], AI that can handle knowledge-based, white-collar work [is] also becoming increasingly competent" (Gershgorn, 2017).

> Consumer AI in the labour market is less apparent as it affects cognitive jobs more than the typically considered manual occupations.

This workplace shift in cognitive jobs is already materialising. For instance, Fukoku Mutual Life Insurance Co. has decided to replace one-third of their payment assessment employees with an AI solution (The Mainichi, 2016). Furthermore, chatbots such as those used in consumer products, are influencing seemingly unrelated areas that have been regarded as inherently human, such as human resources and education (Meister, 2017). Effects are also taking place in creative industries, notwithstanding a recent report by the University of Oxford that lists originality as one of the least automatable skills (Frey and Osborne, 2017). One such example is Jukedeck, which uses AI to compose and adapt professional-quality music, offering consumers personalised audio files. The impact on the music industry extends further to music curation, with platforms such as Spotify and SoundCloud using Consumer AI to displace DJs and radio personalities. Another example is

Google's Deep Dream Generator that uses AI to transform consumers' photos into works of art - a skill previously held by artists alone.

A common solution suggested by individuals consulted for this report was government-directed support for training. However, we challenge this solution as potentially flawed on account of insufficient consideration of Consumer AI's unpredictability.

We recommend that issues of training be captured within the government's existing structural unemployment strategy.

We will illustrate this point by reflecting on one of our semi-structured interviews, in which we discussed a recent article from The Economist. Progress in Consumer AI, demonstrated by computer vision in Snapchat for instance, is leading to breakthroughs that are advancing the AI field more generally. We know that progress in computer vision is being applied to routine cognitive occupations, such as the medical field of radiology, and this is influencing the skills required for these roles (The Economist, 2016b). This loose connection from Snapchat to radiology demonstrates our belief that it is too difficult to anticipate how developments in Consumer AI will impact the labour market. As a result, targeted government training may neglect individuals in fields where changes to job skills are less predictable. Instead, we recommend that issues of training be captured within the government's existing structural unemployment strategy (Department of Work and Pensions, 2015; GOV.UK, n.d.).

We acknowledge that it is problematic to comprehensively consider the unintended consequence of Consumer AI in the labour market, without also considering broader implications of automation. While we have challenged some suggestions of training, we firmly believe that this area warrants extensive future research - taking into account the topics that came up during our conversations: the role of education reform, societal and identity crises, and the responsibilities of businesses.

KEY FINDINGS AND RECOMMENDATIONS

Key Findings and Recommendations

Through analysis of our primary research and relevant secondary sources, we have identified unintended consequences of Consumer AI in ten interest areas. This section provides a summary of our findings. The aim is to encapsulate key insights for each interest area, offering in-context recommendations for policymakers and examples for reference. Each interest area is organised into a table within the two-column matrix below. Through our analysis, we identified the interest areas that are most applicable to organisations trying to understand, build, or take action on Consumer AI in particular. Additionally, policymakers expressed desire to classify unintended consequences, with respect to the character of their effect on consumers. In response, labels have been used to signal both the interest areas that are Most Applicable to Consumer AI and whether the effect on consumers is Direct or Indirect.

1. Consumer Knowledge	
	Most Applicable · Direct
Key Finding	The discrepancy between people's perceived and actual understanding of technology is growing with the deployment of Consumer AI.
Recommendations	We highlight that there is an opportunity for both the private and public sector to provide more information on the workings of AI-reliant products and services.
Examples	Chatbots; video-streaming services; social media.

2. Principal-Agent Relationship	
	Most Applicable · Direct
Key Finding	The traditional principal-agent problem sees a growing prevalence in applications of Consumer AI.
Recommendations	We recommend that governments avoid forcing transparency requirements as this has been shown to be ineffective, and we suspect that it may lead to an unnecessary loss of competitive advantage for the businesses concerned.
Examples	Video-streaming services; ecommerce platforms; navigation services.

3. Agent Scrutiny

Most Applicable | Indirect

Key Finding	Increased use of certain AI techniques leads to the agent scrutiny problem, and thus requires a distinct regulatory approach that does not impede AI research.
Recommendations	We recommend that if governments decide to intervene on matters relating to Consumer AI, such as in the case of hate speech, the focus should be on an agent's inputs and actions rather than the function or program.
Examples	Hate speech articles pushed by curation agents.

4. Discrimination

Most Applicable | Direct

Key Finding	Prejudice in the development and training of intelligent agents can lead to codified discrimination in Consumer AI.
Recommendations	We recommend a review of available options by both government and businesses, to reduce risks of codifying discrimination in implementations of Consumer AI.
Examples	Search results; talent acquisition.

5. Cyber Security

Most Applicable | Indirect

Key Finding	Our research revealed concerns about the heightened impact of data breaches, compromised agents, and offensive capabilities.
Recommendations	We recommend that every organisation consider Consumer AI as a core component of their cyber security strategy.
Examples	Exposure of personal data; compromised chatbots; weaponised malware

6. Market Structure

Most Applicable | Indirect

Key Finding	The market dynamics of Consumer AI may lead to market dominance, however there is no consensus regarding anti-competitive behaviour.
Recommendations	We recommend that policymakers continue studying the issue, and adopt a novel competition lens in the spirit of ensuring that the development of Consumer AI continues to serve consumer interests first - intervening only in cases of anti-competitive behaviour.
Examples	Data regulation as an asset or commodity that can be monopolised.

7. AI Research

Indirect

Key Finding	There is a shift of academics to the private sector and hence the privatisation of future AI research.
Recommendations	We recommend enhanced investment by government to ensure academia remains an attractive environment for researchers to advance the field of AI.
Examples	Yann Lecun to Facebook's AI lab; Fei-Fei Li to Google.

8. AI Developer Ecosystem

Indirect

Key Finding	Collectively: cheaper compute, open data, OSS, and accessible education services, are democratising the AI developer ecosystem.
Recommendations	The only warranted government action is positive incentives that increase engagement in the AI developer ecosystem.
Examples	The open sourcing of TensorFlow, DSSTNE, and Torch.

9. Open Data

Indirect

Key Finding	The implementation of Consumer AI has resulted in an increased demand for quality data, as an integral part of developing and improving AI applications.
Recommendations	Promoting an increase in the availability, usability, and awareness of public data must be a key consideration for policymakers going forward.
Examples	data.gov.uk; data.gov; data.gouv.fr

10. Labour Market

Direct

Key Finding	Consumer AI in the labour market is less apparent as it affects cognitive jobs more than the typically considered manual occupations.
Recommendations	We recommend that issues of training be captured within the government's existing structural unemployment strategy.
Examples	Chatbots; music curators.

CONCLUSION

Conclusion

Artificial intelligence changes how we live and interact in ways not previously imaginable. This sentiment, shared by business leaders, policymakers and academics, from Jeff Bezos, to Barack Obama, and Lord Martin Rees, inspired this inquiry. However, as highlighted, the spread of Consumer AI, from personal assistants, to chatbots, to recommendation engines, and social networks, required us to emphasise the unintended consequences in the sphere of commercial development. By looking at the benefits and drawbacks beyond the intrinsic motivations of implementers, this report identified and explored ten interest areas where Consumer AI impacts us all at business, economic, and social levels.

This report began with an Introduction that explained the recent rise of interest in AI. A special focus on consumer applications led us to the main research question of this report: *What are the unintended consequences of implementing artificial intelligence for personal use?*

This report identified and explored ten interest areas where Consumer AI impacts us all at business, economic, and social levels.

Our Methodology consisted primarily of semi-structured interviews with businesses, policymakers, and academics pertinent to the field of AI. This primary research was combined with a consumer survey and an extensive review of academic literature, secondary sources, and current developments in the field.

Our research led to a working definition of **Consumer AI** as *the commercial development of intelligent agents for personal use.*

Based on our research, we identified Unintended Consequences of Consumer AI, where impacts extend beyond implementers' original intentions, in the following ten interest areas:
(1) Consumer Knowledge, (2) Principal-Agent Relationship, (3) Agent Scrutiny, (4) Discrimination, (5) Cyber Security, (6) Market Structure, (7) AI Research, (8) AI Developer Ecosystem, (9) Open Data, (10) Labour Market.

We provided a thorough analysis in these interest areas by offering background information, discussion of potential benefits and drawbacks, and recommendations to policymakers. These were summarised in the Key Findings and Recommendations. The matrix encapsulated key insights for each interest area, offering in-context recommendations and examples for reference. It also identified the interest areas that are most applicable to Consumer AI in particular. It then classified unintended consequences, with respect to direct or indirect effects on consumers.

An important contribution of this report was the proposal of a working definition for Consumer AI. Applying this definition, we showed how the traditional principal-agent relationship should be adapted to the Consumer AI realm. Further, as this report underlined, some of the identified unintended consequences call for more concrete policymaking, while others simply warrant the continued monitoring of the space. For example, special attention is required for regulatory attempts aimed at achieving greater scrutiny of AI agents. This should focus specifically on an agent's inputs or actions, thereby decoupling it from its function or program. On the other hand, examples of areas within the broader AI field, which also require the attention of policymakers include unintended consequences in the Labour Market and AI Research.

We advise that governments continue to scrutinise the impacts of Consumer AI meticulously, take action when needed, but most importantly provide an environment in which research can continue to take place such that social welfare is maximised. Our hope is that business leaders and policymakers will strive to foster development in the exciting field of Consumer AI, enabling artificial intelligence to enhance productivity, solve complex problems, and improve the lives of those who use it.

Our hope is that business leaders and policymakers will strive to foster development in the exciting field of Consumer AI, enabling artificial intelligence to enhance productivity, solve complex problems, and improve the lives of those who use it.

Our desire is that this report will inspire future research. In particular, more thorough analysis of case studies, impacts, and policy effects of Consumer AI. Additionally, the continuous inquiry into the benefits and drawback of rapidly evolving AI applications will be required.

APPENDIX

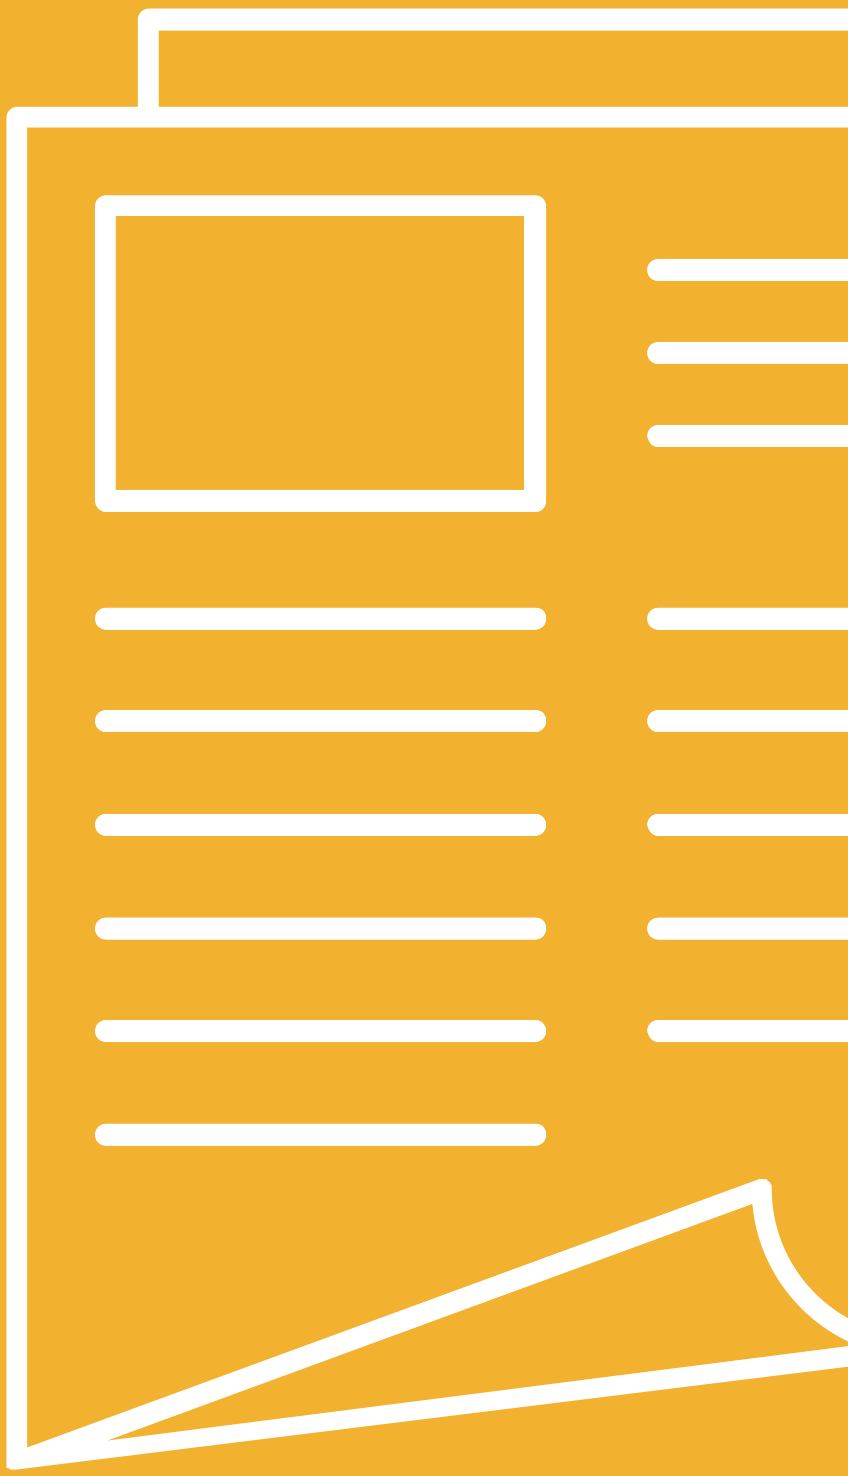

Appendix

Survey Demographics and Summary of Key Findings

Survey Respondents' Age, Gender, and Education

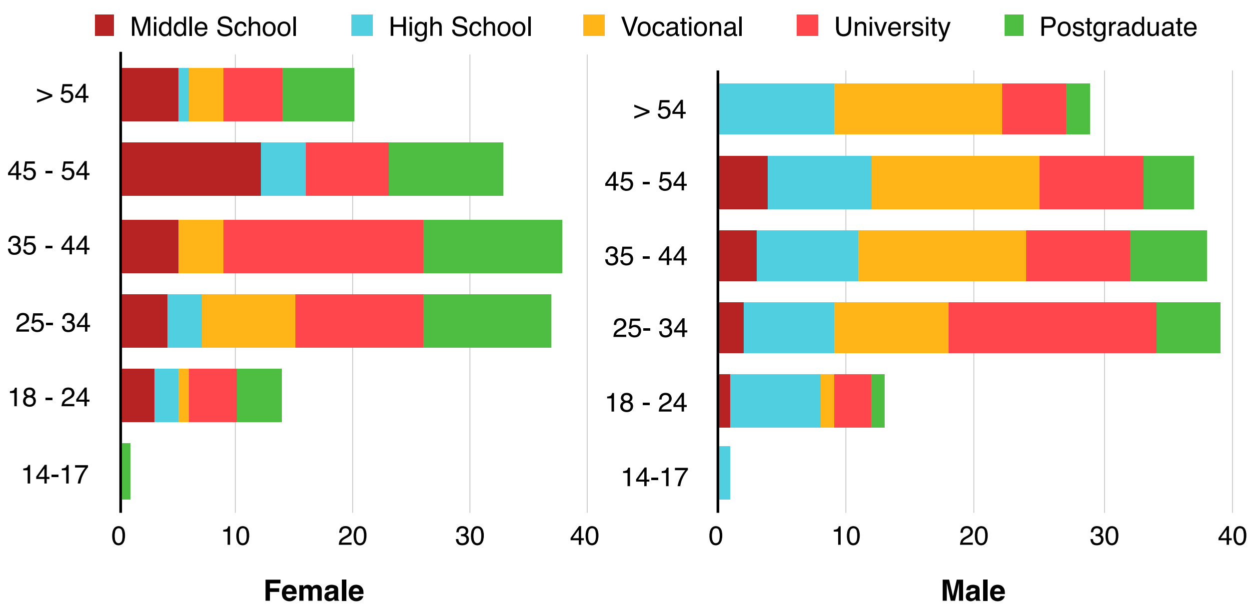

Demonstrated Consumer Knowledge

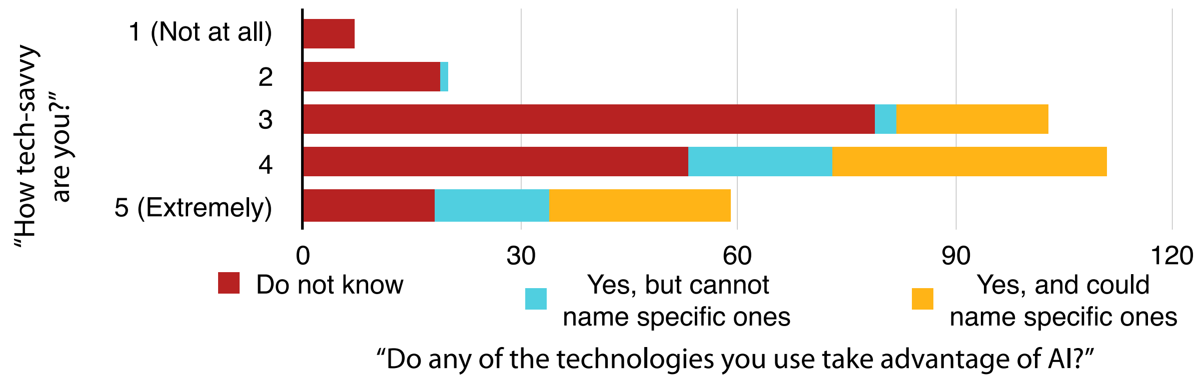

AI's Competitive Advantage

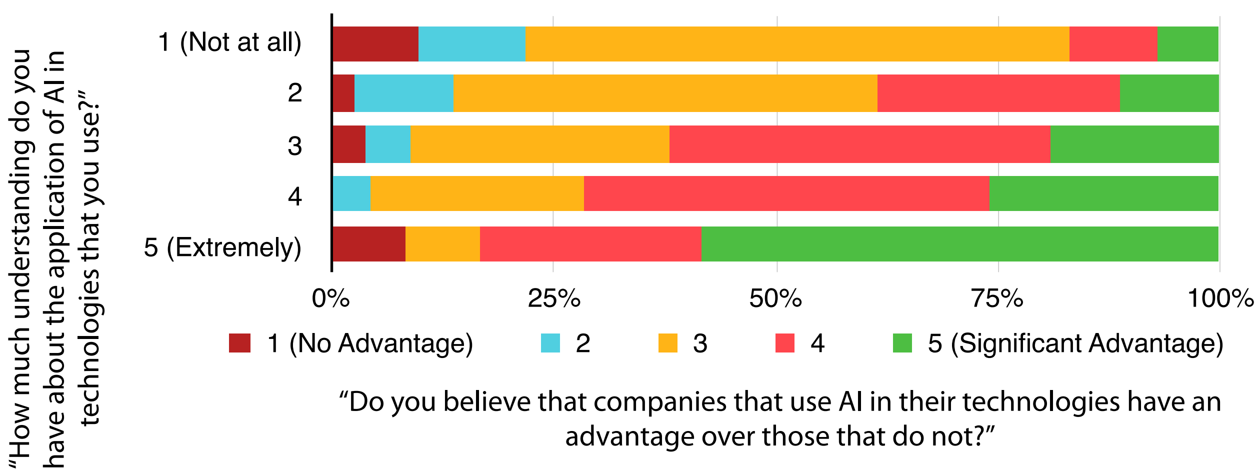

Select Results from Respondents

Which of these do you use on a regular basis?

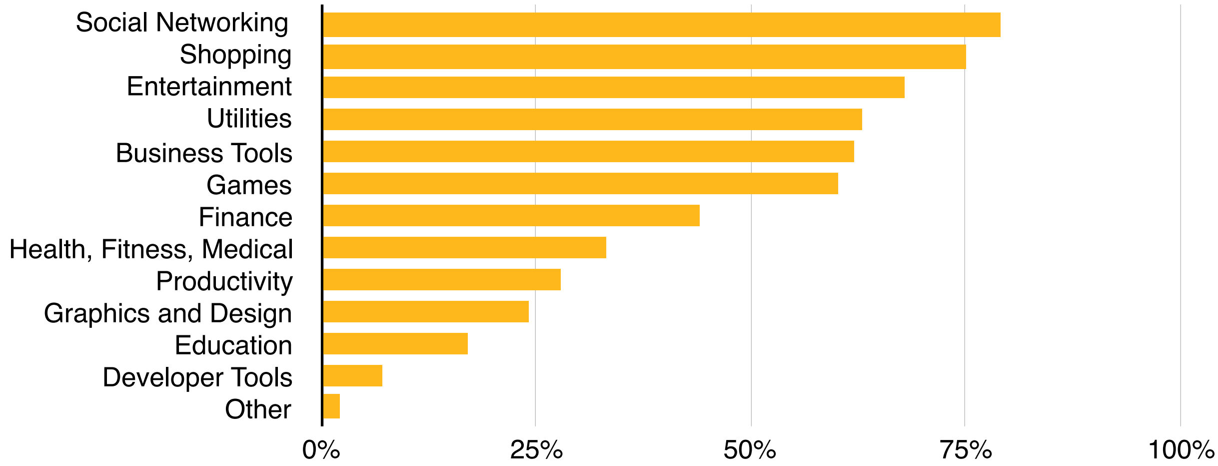

To what extent should the UK government intervene in the use of AI?

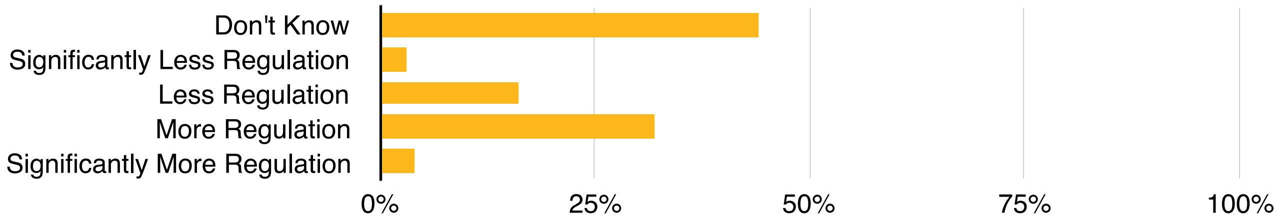

When using new technologies, how concerned are you about the security of the data or information that you share?

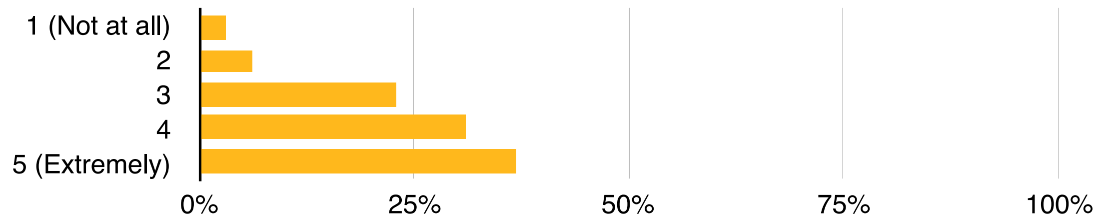

How much information would you like to be provided about how AI technologies work?

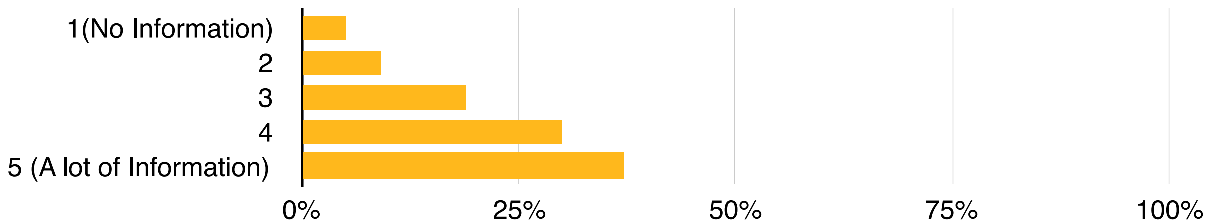

How concerned are you about the use or introduction of AI into technologies that you currently use?

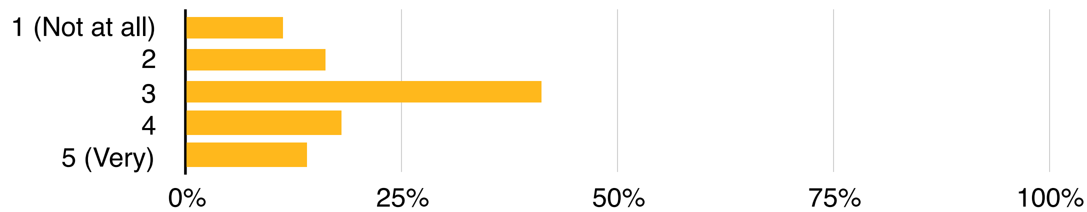

Do you trust technologies that use AI more than those that do not?

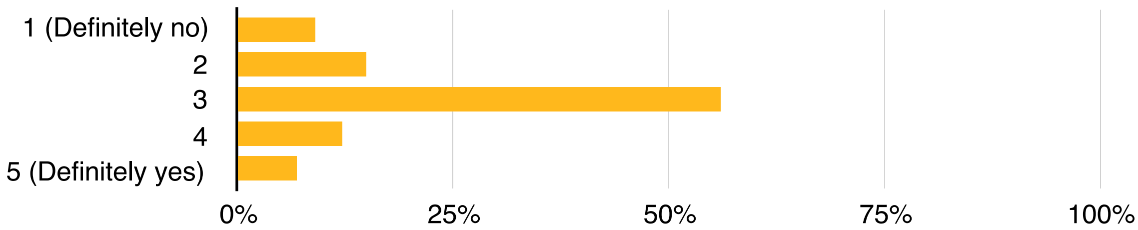

REFERENCE LIST

Reference List

Agrawal, A., Gans, J. and Goldfarb, A. (2016). The simple economics of machine intelligence. Harvard business review. [Online]. Available at: https://hbr.org/2016/11/the-simple-economics-of-machine-intelligence [Accessed: 17 May 2017].

Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J. and Mané, D. (2016). Concrete Problems in AI Safety. [Online]. Available at: http://arxiv.org/abs/1606.06565.

Armstrong, A. (2014). OpenNI To Close. [Online]. Available at: http://www.i-programmer.info/news/194-kinect/7004-openni-to-close-.html [Accessed: 13 May 2017].

Big Innovation Centre. (2017). What is AI? [Online]. Available at: http://biginnovationcentre.com/media/uploads/pdf/APPGAI_Theme%20Report_March%2020_What%20is%20AI.pdf [Accessed: 12 May 2017].

Camp, L. J. and Wolfram, C. (2000). Pricing security. In: Proceedings of the CERT Information Survivability Workshop, 2000, p.31–39. [Online]. Available at: https://pdfs.semanticscholar.org/8385/5bb07d0388ade7e625893a8373adc629983b.pdf [Accessed: 13 May 2017].

Coldewey, D. (2016). Carnegie Mellon's Mayhem AI takes home $2 million from DARPA's Cyber Grand Challenge. TechCrunch. [Online]. Available at: http://social.techcrunch.com/2016/08/05/carnegie-mellons-mayhem-ai-takes-home-2-million-from-darpas-cyber-grand-challenge/ [Accessed: 12 May 2017].

Connolly, K. (2016). Angela Merkel: internet search engines are 'distorting perception'. The Guardian, 27 October. [Online]. Available at: http://www.theguardian.com/world/2016/oct/27/angela-merkel-internet-search-engines-are-distorting-our-perception [Accessed: 12 May 2017].

Constine, J. (2015). Facebook Acquires Wit.ai To Help Its Developers With Speech Recognition And Voice Interfaces. TechCrunch. [Online]. Available at: http://social.techcrunch.com/2015/01/05/facebook-wit-ai/ [Accessed: 13 May 2017].

DARPA. (2016). Cyber Grand Challenge. DARPA. [Online]. Available at: http://archive.darpa.mil/cybergrandchallenge/ [Accessed: 17 May 2017].

data.gov.uk. (n.d.). FAQ | data.gov.uk. [Online]. Available at: https://data.gov.uk/faq [Accessed: 15 May 2017].

Deloitte Analytics. (2012). Open Data: Driving Growth, Ingenuity, and Innovation. Deloitte. [Online]. Available at: https://www2.deloitte.com/content/dam/Deloitte/uk/Documents/deloitte-analytics/open-data-driving-growth-ingenuity-and-innovation.pdf.

DeMers, J. (2016). 7 Technology Trends That Will Dominate 2017. Forbes. [Online]. Available at: http://www.forbes.com/sites/jaysondemers/2016/11/16/7-technology-trends-that-will-dominate-2017/ [Accessed: 15 May 2017].

Department of Work and Pensions. (2015). 2010 to 2015 government policy: employment - GOV.UK. GOV.UK. [Online]. Available at: https://www.gov.uk/government/publications/2010-to-2015-government-policy-employment/2010-to-2015-government-policy-employment [Accessed: 15 May 2017].

Digital Catapult. (2017). Digital Catapult. Digital Catapult. [Online]. Available at: https://digital.catapult.org.uk/about/ [Accessed: 16 May 2017].

Doraiswamy, V. (2017). How AI is detecting unconscious bias in recruitment - Digital Catapult Centre. Digital Catapult Centre. [Online]. Available at: https://www.digitalcatapultcentre.org.uk/ai-detecting-unconscious-bias-recruitment/ [Accessed: 17 May 2017].

Eisenhardt, K. M. (1989). Agency Theory: An Assessment and Review. Academy of management review. Academy of Management, 14 (1), p.57–74. [Online]. Available at: doi:10.5465/AMR.1989.4279003.

European Union. (2016). REGULATION (EU) 2016/ 679 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL. [Online]. Available at: http://ec.europa.eu/justice/data-protection/reform/files/regulation_oj_en.pdf [Accessed: 13 May 2017].

Executive Office of the President of the United States of America. (2017). Artificial Intelligence, Automation, and the Economy. CreateSpace Independent Publishing Platform. [Online]. Available at: https://www.whitehouse.gov/sites/whitehouse.gov/files/images/EMBARGOED%20AI%20Economy%20Report.pdf [Accessed: 12 May 2017].

Farrell, J. and Klemperer, P. (2006). Coordination and Lock-In: Competition with Switching Costs and Network Effects. [Online]. Available at: http://www.escholarship.org/uc/item/9n26k7v1 [Accessed: 13 May 2017].

Feduzi, A. and Runde, J. (2014). Uncovering unknown unknowns: Towards a Baconian approach to management decision-making. Organizational behavior and human decision processes, 124 (2), p.268–283. [Online]. Available at: http://www.sciencedirect.com/science/article/pii/S0749597814000302.

Finley, K. (2016a). Amazon's Giving Away the AI Behind Its Product Recommendations. Wired, 16 May. [Online]. Available at: https://www.wired.com/2016/05/amazons-giving-away-ai-behind-product-recommendations/ [Accessed: 14 May 2017].

Finley, K. (2016b). Google Is About to Supercharge Its TensorFlow Open Source AI. Wired. [Online]. Available at: https://www.wired.com/2016/04/google-supercharging-tensorflow-open-source-ai/ [Accessed: 14 May 2017].

Frey, C. B. and Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? Technological forecasting and social change, 114, p.254–280. [Online]. Available at: http://www.sciencedirect.com/science/article/pii/S0040162516302244.

Furman, J. (2016). Is This Time Different? The Opportunities and Challenges of Artificial Intelligence. [Online]. Available at: http://sites.nationalacademies.org/cs/groups/pgasite/documents/webpage/pga_177042.pdf.

Gallaugher, J. M. and Wang, Y.-M. (2002). Understanding Network Effects in Software Markets: Evidence from Web Server Pricing. The Mississippi quarterly, 26 (4), Management Information Systems Research Center, University of Minnesota, p.303–327. [Online]. Available at: doi:10.2307/4132311 [Accessed: 13 May 2017].

Garcia, M. (2017). How to Keep Your AI From Turning Into a Racist Monster. Wired. [Online]. Available at: https://www.wired.com/2017/02/keep-ai-turning-racist-monster/ [Accessed: 17 May 2017].

Garside, J. (2016). Labour calls for closer scrutiny of tech firms and their algorithms. The Guardian, 18 December. [Online]. Available at: http://www.theguardian.com/business/2016/dec/18/labour-calls-for-regulation-of-algorithms-used-by-tech-firms [Accessed: 15 May 2017].

Gershenson, C. (2003). Artificial Neural Networks for Beginners. Datajobs. [Online]. Available at: https://datajobs.com/data-science-repo/Neural-Net-[Carlos-Gershenson].pdf [Accessed: 14 May 2017].

Gershgorn, D. (2017). Japanese white-collar workers are already being replaced by artificial intelligence. Quartz, Quartz. [Online]. Available at: https://qz.com/875491/japanese-white-collar-workers-are-already-being-replaced-by-artificial-intelligence/ [Accessed: 14 May 2017].

Gong, G., Qiang, S. and Wang, J. (2009). Information security measures and regulation research. In: 2009 International Conference on Management Science and Engineering, 2009, p.2184–2189. [Online]. Available at: doi:10.1109/ICMSE.2009.5317613.

Gordon, L. A., Loeb, M. P., Lucyshyn, W. and Zhou, L. (2015). Externalities and the Magnitude of Cyber Security Underinvestment by Private Sector Firms: A Modification of the Gordon-Loeb Model. Journal of Information Security, 06 (01), p.24–30. [Online]. Available at: doi:10.4236/jis.2015.61003.

GOV.UK. (n.d.). Employment - Policy areas - GOV.UK. GOV.UK. [Online]. Available at: https://www.gov.uk/government/topics/employment [Accessed: 15 May 2017].

Greenberg, A. (2016). Obama's Concerned an AI Could Hack America's Nukes. Wired, 12 October. [Online]. Available at: https://www.wired.com/2016/10/obamas-concerned-ai-hack-americas-nukes/ [Accessed: 13 May 2017].

Hardt, M., Price, E. and Srebro, N. (2016). Equality of Opportunity in Supervised Learning. [Online]. Available at: https://research.google.com/bigpicture/attacking-discrimination-in-ml/ [Accessed: 15 May 2017].

Harris, M. and Raviv, A. (1978). Some results on incentive contracts with applications to education and employment, health insurance, and law enforcement. The American economic review, 68 (1), p.20–30. [Online]. Available at: https://www.ncbi.nlm.nih.gov/pubmed/10307940.

Hassabis, D. (2017). The mind in the machine: Demis Hassabis on artificial intelligence. Financial Times. [Online]. Available at: https://www.ft.com/content/048f418c-2487-11e7-a34a-538b4cb30025 [Accessed: 17 May 2017].

HM Government. (2016). UK Open Government Action Plan 2016-2018. GOV.UK. [Online]. Available at: https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/522781/UK_Open_Government_National_Action_Plan_2016-18.pdf.

IBM. (2017). IBM - What is big data? IBM. [Online]. Available at: https://www-01.ibm.com/software/data/bigdata/what-is-big-data.html [Accessed: 16 May 2017].

Joblint. (n.d.). Joblint. Joblint. [Online]. Available at: http://joblint.org/ [Accessed: 12 May 2017].

Kagel, J. (2013). Facebook Hires 'Deep Learning' Expert To Lead Artificial Intelligence Research Lab. Fast Company, Fast Company. [Online]. Available at: https://www.fastcompany.com/3023169/facebook-hires-deep-learning-expert-to-lead-artificial-intelligence-research-lab [Accessed: 13 May 2017].

Kelly, K. (2014). The Three Breakthroughs That Have Finally Unleashed AI on the World. Wired. [Online]. Available at: https://www.wired.com/2014/10/future-of-artificial-intelligence/ [Accessed: 16 May 2017].

Kelnar, D. (2016). Artificial Intelligence in the UK: Landscape and learnings from 226 startups. Medium, MMC writes. [Online]. Available at: https://medium.com/mmc-writes/artificial-intelligence-in-the-uk-landscape-and-learnings-from-226-startups-70b9551f3e4c [Accessed: 13 May 2017].

Kemerer, C. F., Liu, C. Z. and Smith, M. D. (2013). Strategies for Tomorrow's 'WinnersTake-Some' Digital Goods Markets. [Online]. Available at: doi:10.1145/2447976.2447994 [Accessed: 17 May 2017].

Knight, W. (2015). The Roomba Now Sees and Maps a Home. Technology Review. [Online]. Available at: https://www.technologyreview.com/s/541326/the-roomba-now-sees-and-maps-a-home/ [Accessed: 14 May 2017].

Knight, W. (2017). The Dark Secret at the Heart of AI. Technology Review. [Online]. Available at: https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/ [Accessed: 14 May 2017].

Kunreuther, H. and Heal, G. (2003). Interdependent Security. Journal of risk and uncertainty, 26 (2-3), Kluwer Academic Publishers, p.231–249. [Online]. Available at: doi:10.1023/A:1024119208153 [Accessed: 17 May 2017].

Laffont, J.-J. and Martimort, D. (2009). The Theory of Incentives: The Principal-Agent Model. Princeton University Press. [Online]. Available at: https://market.android.com/details?id=book-Yf1TwtIuNf8C.

Lardinois, F., Lynley, M. and Mannes, J. (2017). Google is acquiring data science community Kaggle. TechCrunch. [Online]. Available at: http://social.techcrunch.com/2017/03/07/google-is-acquiring-data-science-community-kaggle/ [Accessed: 13 May 2017].

Lee, C.-Y., Xie, S., Gallagher, P., Zhang, Z. and Tu, Z. (2015). Deeply-Supervised Nets. In: Artificial Intelligence and Statistics, 21 February 2015, p.562–570. [Online]. Available at: http://proceedings.mlr.press/v38/lee15a.html [Accessed: 17 May 2017].

Li, H. (2011). Capital structure on agency costs in Chinese listed firms. International Journal of Business Governance & Ethics, 1 (2), p.295–308.

Lunden, I. (2016). Twitter pays up to $150M for Magic Pony Technology, which uses neural networks to improve images. TechCrunch. [Online]. Available at: http://social.techcrunch.com/2016/06/20/twitter-is-buying-magic-pony-technology-which-uses-neural-networks-to-improve-images/ [Accessed: 14 May 2017].

Maher, B. (2016). Can a video game company tame toxic behaviour? Nature, 531 (7596), p. 568–571. [Online]. Available at: doi:10.1038/531568a.

Mannes, J. (2017). Galvanize will teach students how to use IBM Watson APIs with new machine learning course. TechCrunch. [Online]. Available at: http://social.techcrunch.com/2017/03/19/galvanize-will-teach-students-how-to-use-ibm-watson-apis-with-new-machine-learning-course/ [Accessed: 14 May 2017].

Marc'aurelio, R. and Lecun, Y. (2013). Deep Learning Tutorial. NYU. [Online]. Available at: http://www.cs.nyu.edu/~yann/talks/lecun-ranzato-icml2013.pdf.

Martin, W., Fernanda, V. and Moritz, H. (2016). Attack discrimination with smarter machine learning. Research Google. [Online]. Available at: https://research.google.com/bigpicture/attacking-discrimination-in-ml/ [Accessed: 15 May 2017].

Meister, J. (2017). The Future Of Work: The Intersection Of Artificial Intelligence And Human Resources. Forbes. [Online]. Available at: http://www.forbes.com/sites/jeannemeister/2017/03/01/the-future-of-work-the-intersection-of-artificial-intelligence-and-human-resources/ [Accessed: 14 May 2017].

Merton, R. K. (1936). The Unanticipated Consequences of Purposive Social Action. American sociological review, 1 (6), [American Sociological Association, Sage Publications, Inc.], p.894–904. [Online]. Available at: doi:10.2307/2084615.

Moore, D. A. and Healy, P. J. (2008). The trouble with overconfidence. Psychological review, 115 (2), p.502–517. [Online]. Available at: doi:10.1037/0033-295X.115.2.502.

Munday, S. C. R. (2000). Markets and Market Failure. Heinemann.

Murgia, M. (2017). UK government looks for opportunities in AI and robotics. Financial Times. [Online]. Available at: https://www.ft.com/content/9c1c80b6-fcbc-11e6-96f8-3700c5664d30 [Accessed: 13 May 2017].

Newman, D. (2017). 17 Top Enterprise Tech Trends For 2017 And This Year's CES. Forbes. [Online]. Available at: http://www.forbes.com/sites/danielnewman/2017/01/03/17-top-enterprise-tech-trends-for-2017-and-this-years-ces/ [Accessed: 15 May 2017].

Nikkinen, J. and Sahlström, P. (2004). Does Agency Theory Provide a General Framework for Audit Pricing? International Journal of Accounting, Auditing and Peformance Evaluation, 8 (3), Blackwell Publishing Ltd., p.253–262. [Online]. Available at: doi: 10.1111/j.1099-1123.2004.00094.x.

Norton, R. (2002). Unintended Consequences: The Concise Encyclopedia of Economics | Library of Economics and Liberty. [Online]. Available at: http://www.econlib.org/library/Enc/UnintendedConsequences.html [Accessed: 12 May 2017].

Novet, J. (2016). Amazon open-sources its own deep learning software, DSSTNE. VentureBeat. [Online]. Available at: https://venturebeat.com/2016/05/11/amazon-open-sources-its-own-deep-learning-software-dsstne/ [Accessed: 14 May 2017].

O2. (2015). Customer data – an asset or a liability? O2. [Online]. Available at: http://news.o2.co.uk/2015/10/16/customer-data-asset-liability/ [Accessed: 17 May 2017].

Open Data Barometer. (2015). ODB Global Report Third Edition. Open Data Barometer. [Online]. Available at: http://opendatabarometer.org/doc/3rdEdition/ODB-3rdEdition-GlobalReport.pdf.

Open Data Institute. (2017). Research: Open data means business. Open Data Institute. [Online]. Available at: https://theodi.org/open-data-means-business [Accessed: 15 May 2017].

Panetta, K. (2016). Gartner's Top 10 Strategic Technology Trends for 2017 - Smarter With Gartner. Smarter With Gartner. [Online]. Available at: http://www.gartner.com/smarterwithgartner/gartners-top-10-technology-trends-2017/ [Accessed: 15 May 2017].

Perens, B. and Others. (1999). The open source definition. OpenSource.org, Sebastopol, CA: O'Reilly. [Online]. Available at: https://opensource.org/osd-annotated [Accessed: 13 May 2017].

Ponemon Institute. (2011). Reputation Impact of a Data Breach. Experian. [Online]. Available at: https://www.experian.com/assets/data-breach/white-papers/reputation-study.pdf [Accessed: 17 May 2017].

Ponemon Institute. (2012). Aftermath of a Data Breach Study. Experian. [Online]. Available at: http://www.ponemon.org/local/upload/file/Aftermath_of_a_Data_Breach_WP_Final%20.pdf.

Ponemon Institute. (2016). 2016 Cost of Data Breach Study: Global Analysis. [Online]. Available at: https://app.clickdimensions.com/blob/softchoicecom-anjf0/files/ponemon.pdf.

Press, G. (2016a). A Very Short History Of Artificial Intelligence (AI). Forbes. [Online]. Available at: http://www.forbes.com/sites/gilpress/2016/12/30/a-very-short-history-of-artificial-intelligence-ai/ [Accessed: 13 May 2017].

Press, G. (2016b). Forrester Predicts Investment In Artificial Intelligence Will Grow 300% in 2017. Forbes. [Online]. Available at: http://www.forbes.com/sites/gilpress/2016/11/01/forrester-predicts-investment-in-artificial-intelligence-will-grow-300-in-2017/ [Accessed: 14 May 2017].

Price, R. (2016). Microsoft is deleting its AI chatbot's incredibly racist tweets. Business Insider. [Online]. Available at: http://uk.businessinsider.com/microsoft-deletes-racist-genocidal-tweets-from-ai-chatbot-tay-2016-3 [Accessed: 12 May 2017].

Pringle, R., Michael, K. and Michael, M. (2016). Unintended Consequences of Living with AI: The Paradox of Technological Potential?Part II [Guest Editorial]. IEEE Technology and Society Magazine, 35 (4), p.17–21. [Online]. Available at: doi:10.1109/MTS.2016.2632978.

Rao, L. (2017). Apple Just Acquired This Little-Known Artificial Intelligence Startup. Fortune, Fortune. [Online]. Available at: http://fortune.com/2017/05/13/apple-lattice/ [Accessed: 14 May 2017].

Royal Society. (2017). Machine learning: the power and promise of computers that learn by example. Royal Society. [Online]. Available at: https://royalsociety.org/~/media/policy/projects/machine-learning/publications/machine-learning-report.pdf.

Russell, S. J. and Norvig, P. (2010). Artificial Intelligence: A Modern Approach. Prentice Hall. [Online]. Available at: https://market.android.com/details?id=book-8jZBksh-bUMC.

Salanié, B. (2000). Microeconomics of Market Failures. MIT Press.

Sato, K. (2016). How a Japanese cucumber farmer is using deep learning and TensorFlow. Google Cloud Platform. [Online]. Available at: https://cloud.google.com/blog/big-data/2016/08/how-a-japanese-cucumber-farmer-is-using-deep-learning-and-tensorflow [Accessed: 14 May 2017].

Shadbolt, N. (2015). Comment: Building the UK's National Data Infrastructure. Financial Times. [Online]. Available at: https://www.ft.com/content/ef4730f0-10ff-30ce-b05e-74918dc48eac [Accessed: 14 May 2017].

Simonite, T. (2017). Nvidia CEO: Software Is Eating the World, but AI Is Going to Eat Software. Technology Review. [Online]. Available at: https://www.technologyreview.com/s/607831/nvidia-ceo-software-is-eating-the-world-but-ai-is-going-to-eat-software/ [Accessed: 16 May 2017].

Smith, M. (2016). Public Input and Next Steps on the Future of Artificial Intelligence. Medium. [Online]. Available at: https://medium.com/@USCTO/public-input- and-next-steps-on-the-future-of-artificialintelligence-458b82059fc3 [Accessed: 12 May 2017].

Stacy Jones, J. T. (2015). Tech company diversity: How do the big players compare? Fortune, Fortune. [Online]. Available at: http://fortune.com/2015/07/30/tech-companies-diveristy/ [Accessed: 13 May 2017].

Summers, L. (2016). Data collection is the ultimate public good. Financial Times. [Online]. Available at: https://www.ft.com/content/6d3d019e-9a31-33b7-a7ec-e79f4f108aa0 [Accessed: 15 May 2017].

Symantec. (n.d.). Endpoint Protection – Machine Learning Security | Symantec. www.symantec.com. [Online]. Available at: https://www.symantec.com/products/endpoint-hybrid-cloud-security/endpoint/endpoint-protection [Accessed: 17 May 2017].

The Economist. (2015). The dawn of artificial intelligence. The Economist, The Economist. [Online]. Available at: http://www.economist.com/news/leaders/21650543-powerful-computers-will-reshape-humanitys-future-how-ensure-promise-outweighs [Accessed: 15 May 2017].

The Economist. (2016a). After Moore's law | Technology Quarterly. The Economist. [Online]. Available at: http://www.economist.com/technology-quarterly/2016-03-12/after-moores-law [Accessed: 14 May 2017].

The Economist. (2016b). Automation and anxiety. The Economist, The Economist. [Online]. Available at: http://www.economist.com/news/special-report/21700758-will-smarter-machines-cause-mass-unemployment-automation-and-anxiety [Accessed: 14 May 2017].

The Economist. (2016c). Million-dollar babies. The Economist, The Economist. [Online]. Available at: http://www.economist.com/news/business/21695908-silicon-valley-fights-talent-universities-struggle-hold-their [Accessed: 13 May 2017].

The Economist. (2016d). What DeepMind brings to Alphabet. The Economist, The Economist. [Online]. Available at: http://www.economist.com/news/business/21711946-ai-firms-main-value-alphabet-new-kind-algorithm-factory-what-deepmind-brings [Accessed: 14 May 2017].

The Economist. (2017a). Data is giving rise to a new economy. The Economist, The Economist. [Online]. Available at: http://www.economist.com/news/briefing/21721634-how-it-shaping-up-data-giving-rise-new-economy [Accessed: 14 May 2017].

The Economist. (2017b). The world's most valuable resource is no longer oil, but data. The Economist, The Economist. [Online]. Available at: http://www.economist.com/news/leaders/21721656-data-economy-demands-new-approach-antitrust-rules-worlds-most-valuable-resource [Accessed: 13 May 2017].

The Mainichi. (2016). Insurance firm to replace human workers with AI system - The Mainichi. The Mainichi. [Online]. Available at: http://mainichi.jp/english/articles/20161230/p2a/00m/0na/005000c [Accessed: 14 May 2017].

UK Cabinet Office. (2012). Open Data White Paper. GOV.UK. [Online]. Available at: https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/78946/CM8353_acc.pdf.

UK Government. (1998). Data Protection Act 1998. Statute Law Database. [Online]. Available at: http://www.legislation.gov.uk/ukpga/1998/29/contents [Accessed: 13 May 2017].

Vanian, J. (2016). Google Hires Two Artificial Intelligence Experts To Lead Machine Learning Team. Fortune, Fortune. [Online]. Available at: http://fortune.com/2016/11/15/google-fei-fei-li-snapchat-machine-learning/ [Accessed: 13 May 2017].

Wachter, S., Mittelstadt, B. and Floridi, L. (2016). Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation. [Online]. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2903469 [Accessed: 15 May 2017].

Waters, R. (2016). AI academic warns on brain drain to tech groups. Financial Times. [Online]. Available at: https://www.ft.com/content/298e2ac0-b010-11e6-a37c-f4a01f1b0fa1 [Accessed: 15 May 2017].

Weber Shandwick. (2016). AI-READY OR NOT: ARTIFICIAL INTELLIGENCE HERE WE COME! Weber Shandwick. [Online]. Available at: http://www.webershandwick.com/uploads/news/files/AI-Ready-or-Not-report-Oct12-FINAL.pdf.

Weinberger, D. (2016). How the father of the World Wide Web plans to reclaim it from Facebook and Google. Digital Trends. [Online]. Available at: https://www.digitaltrends.com/web/ways-to-decentralize-the-web/ [Accessed: 14 May 2017].

West, D. M. (2015). What happens if robots take the jobs? The impact of emerging technologies on employment and public policy. Centre for Technology Innovation at Brookings, insidepolitics.org. [Online]. Available at: http://www.insidepolitics.org/brookingsreports/robots.pdf.

Wharton. (2017). The Democratization of Machine Learning: What It Means for Tech Innovation. Knowledge@Wharton. [Online]. Available at: http://knowledge.wharton.upenn.edu/article/democratization-ai-means-tech-innovation/ [Accessed: 14 May 2017].

World White Web. (2016). World White Web. World White Web. [Online]. Available at: http://www.worldwhiteweb.net/ [Accessed: 13 May 2017].

Yao, M. (2017). Why We Need To Democratize Artificial Intelligence Education. Forbes. [Online]. Available at: http://www.forbes.com/sites/mariyayao/2017/04/10/why-we-need-to-democratize-ai-machine-learning-education/ [Accessed: 13 May 2017].

Zhang, M. (2015). Google Photos Tags Two African-Americans As Gorillas Through Facial Recognition Software. Forbes. [Online]. Available at: http://www.forbes.com/sites/mzhang/2015/07/01/google-photos-tags-two-african-americans-as-gorillas-through-facial-recognition-software/ [Accessed: 12 May 2017].

UNIVERSITY OF CAMBRIDGE
Judge Business School

#cvr #bmd #ai4lyf