



International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: <https://www.svedbergopen.com/>



Research Paper

Open Access

The Future of Artificial Intelligence

Andrew C. Scott¹, José R. Solórzano², Jonathan D. Moyer³, and Barry B. Hughes^{4*}

¹Frederick S. Pardee Center for International Futures, Josef Korbel School of International Studies, University of Denver, 2201 S Gaylord St #120, Denver, CO 80208, United States. E-mail: acs.14.cfc@gmail.com

²Frederick S. Pardee Center for International Futures, Josef Korbel School of International Studies, University of Denver, 2201 S Gaylord St #120, Denver, CO 80208, United States. E-mail: jsolorza@du.edu

³Frederick S. Pardee Center for International Futures, Josef Korbel School of International Studies, University of Denver, 2201 S Gaylord St #120, Denver, CO 80208, United States. E-mail: jmoyer@du.edu

⁴Frederick S. Pardee Center for International Futures, Josef Korbel School of International Studies, University of Denver, 2201 S Gaylord St #120, Denver, CO 80208, United States. E-mail: bhughes@du.edu

Article Info

Volume 2, Issue 1, January 2022

Received : 15 December 2021

Accepted : 06 January 2022

Published : 18 January 2022

[doi: 10.51483/IJAIML.2.1.2022.1-37](https://doi.org/10.51483/IJAIML.2.1.2022.1-37)

Abstract

We present a global Artificial Intelligence (AI) conceptual framework, operationalization, and forecast to the year 2100. A series of AI indices were developed within the International Futures (IFs) integrated assessment platform, a quantitative macro-level system that produces dynamic forecasts for 186 countries. IFs models extensively interconnected aspects of global human development, including: agriculture, economics, demographics, energy, infrastructure, environment, water, governance, health, education, finance, technology, and international politics. We conceptualize AI in three categories: narrow AI, general artificial intelligence (AGI), and superintelligence. Today's AI consists of six basic and narrow AI technologies: computer vision, machine learning, natural language processing, the Internet of Things (IoT), robotics, and reasoning. As an index score for all approaches 10, we forecast AGI technology to become available, representing it with a machine IQ index score, roughly analogous to human IQ scores. The emergence of AGI is constrained by the rate of improvement in and development of machine reasoning and associated technologies. When machine IQ scores approach superhuman levels, we forecast the emergence of superintelligent AI. The current path forecast estimates that AGI could appear between 2040 and 2050. Superintelligent AI is forecast to be developed close to the end of the current century. We frame the current path with faster and slower scenarios of development and facilitate analysis of alternative scenarios. Future work can assess the complex impacts of AI development on human society, including economic productivity, labor, international trade, and energy systems.

Keywords: *Artificial Intelligence, International Futures (IFs), Superintelligence, Modeling, Forecasting*

© 2022 Andrew C. Scott et al. This is an open access article under the CC BY license (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

* Corresponding author: Barry B. Hughes, Frederick S. Pardee Center for International Futures, Josef Korbel School of International Studies, University of Denver, 2201 S Gaylord St #120, Denver, CO 80208, United States. E-mail: bhughes@du.edu

1. Introduction and Overview¹

The term Artificial Intelligence, or AI, conjures widely different images and expectations for many different people. Some imagine a world filled by autonomous vehicles zipping around without human input. Others may imagine a world where intelligent robots work alongside humans helping to remove much of the drudgery and daily toil from their lives. Some see rapid advances in healthcare and healthcare technologies, enabling humans to live healthier, fitter, and longer lives. Some may see a world where AI becomes the great equalizer, lowering the cost of production and making a wide range of goods available to broad swathes of the population. And yet for some, AI conjures fear and foreboding, a world characterized by mass dislocation of labor and inequality, generating vast social instability. The great fear is that AI comes to surpass human capability with devastating and unknown consequences.

Despite these widely different predictions of future AI and human interaction, AI technologies today remain remarkably limited and narrow, capable of generating only simple outputs like responding to questions, or identifying specific objects within images, or identifying anomalies from complex patterns of data. The world of autonomous agents with intelligence equaling or even exceeding that of humans is still largely a fantasy. And yet today's narrow AI technologies are advancing rapidly, doubling or even tripling their performance over the past five to ten years. AI has been called the "Fourth Industrial Revolution," (Schwab and Samans, 2016a) a recognition its potential impact across a number of important sectors of human development.

AI will have far-reaching effects on the economy; enhancing productivity while at the same time shifting the value-add away from labor and towards capital-intensive machinery and industries. The direct effects on labor are hotly debated. AI technologies are already replacing labor in manufacturing and in some service sectors today, and pessimists suggest this is a harbinger of a broader trend that will lead to massive hollowing out of jobs brought on by automation of tasks and employment. Optimists counter this by pointing out that technology has historically been a net job creator, leading to the development of entirely new industries and specializations previously unavailable. AI will simply free up human capital to pursue more productive and meaningful pursuits, they say. In other sectors, the impact will be similarly broad. Autonomous vehicles could fundamentally restructure transportation infrastructure, reduce traffic accidents and associated congestion. AI could help drive renewable energy generation and improve demand-side efficiencies, leading to a massive growth in renewable power. AI could personalize education service delivery and produce tools that allow for life-long learning. AI's potential is both wide and deep and only beginning to be realized.

Given AI's rapid advance and associated consequences, a model of AI development with the capacity for scenario analysis to explore forward impacts is valuable. The purpose of this paper is to document an effort to build a quantitative forecast of AI within the IFs integrated assessment platform, housed at the Frederick S Pardee Center for International Futures. While no modeling effort can fully capture the diverse impacts of the AI revolution, the integrated nature of the IFs system leaves it uniquely placed to model AI and explore the forward impacts. The AI representation is designed to be uniquely customizable within IFs allowing users to calibrate the representation based on their own conceptions of how the field is progressing.

We begin with consideration of some of the drivers of AI development, in particular: hardware and software development, the rise of Big Data and cloud computing, information and communication technology penetration rates, and growing investment. We discuss the construction of the indices and initial model results, and then suggest some potential sectors to explore the impact of AI within the IFs framework in future research. We highlight the potential impact on economic productivity, labor, and global trade patterns, particularly within the context of greater capacity for localized production and renewable energy generation.

2. Conceptualizing the Field of AI

AI refers generally to the development of machines and autonomous agents able to perform tasks normally requiring human-level intelligence. The field of AI was formally identified in the 1950s, and subsequent development was uneven, punctuated by prolonged periods of reduced attention and funding. Over the past five to ten years there has been renewed interest, particularly from commercial entities, coupled with rapid investment in AI and AI-related technologies. By one estimate, in 2015 technology companies spent close to \$8.5 bn on deals and investments in AI, four times as much as 2010 (*The Economist*, 2016). In 2014 and 2015 alone, eight global technology firms (including major firms like Google and Microsoft) made 26 acquisitions of start-ups producing AI technologies for

¹ This article has largely the same content as a paper developed earlier that is available on SSRN as "Modeling Artificial Intelligence and Exploring its Impact."

an estimated \$5 bn (Chen et al., 2016). In February 2017 Ford motor company announced it would invest \$1 bn into technologies to promote research on self-driving cars (Isaac and Boudette, 2017). These same technology giants and industry investors are currently engaged in a fierce competition for talent to develop an AI platform that will become industry standard, allowing that company, or set of companies, to control its development for years to come.

The field of AI is changing rapidly; it is something of a “Wild Wild West” for both research and investment. The 2016 Association for the Advancement of Artificial Intelligence Conference, one of the largest, accepted submissions to over 30 sub-disciplines of AI. Between 2012 and 2015, the *Wall Street Journal* estimated that close to 170 startups opened in Silicon Valley that were focused on AI (Waters, 2015). To help conceptualize such a large and varied field, we have drawn on multiple threads of research to build a representation in IFs that proceeds along three major categories or typologies: narrow, general, and super AI.

2.1. Major AI Typologies

Narrow (Weak) AI: Refers to specialized systems designed to perform only one task, such as speech and image recognition, or machine translation. Almost all recent progress in the field is happening within the confines of the narrow AI. Examples of narrow AI include: Apple iPhone’s intelligent personal assistant Siri, Alexa from Amazon echo, Google’s automated translation feature, video game AI, and automated customer support. Narrow AI’s rapid growth and development is being driven by improving technology, rising investment, and a growing recognition of the substantial commercial and social benefits accruing from these technologies.

General (Strong) AI (AGI): Seeks to create a single system that exhibits general human intelligence across any cognitive area including language, perception, reasoning, creativity, and planning. Constructing machines with AGI is extremely complex, and they have yet to be created. While the development of AGI may have been one of the original goals of the AI movement, there is a large amount of uncertainty around when AGI will emerge. Most research in recent years has not focused on AGI and there is no comprehensive roadmap toward such an outcome (Stone et al., 2016).

Superintelligent AI: AI superintelligence refers to an intellect “any intellect that greatly exceeds the cognitive performance of humans in virtually all domains of interest” (Bostrom, 2014, p. 26). This broad definition does not classify what form superintelligence could take, whether a network of computers, a robot, or something else entirely. It also treats superintelligence as a monolithic entity, when in fact it may be possible to create machines with “superabilities,” which we currently lack the ability to define and measure (Hernández-Orallo, 2017, p. 24). Researchers have suggested that the advent of AGI will create a positive feedback loop in both research and investment, leading to the development of superintelligent machines (Bostrom, 1998).

3. A Survey of Drivers of AI

To understand and identify trends in AI development, a survey of the key conceptual and technical drivers is important. Important drivers include: hardware and software development, commercial investment, Big Data and cloud computing, and levels of Information and Communication Technology (ICT) penetration. We recognize that this list may not be comprehensive nor exhaustive but believe that these areas represent important proximate drivers of AI and important conceptual building blocks of the AI forecasting capability in IFs.

3.1. Hardware Development

AI development relies on two major technological thrusts: hardware and software. Hardware, or computing and processing power, has traditionally been conceived in relation to Moore’s Law. Named for Intel co-founder Gordon Moore, it refers to his observation in 1965 that the number of transistors on a computing microchip had doubled every year since their intervention, and was forecast to continue along that trajectory (Figure 1).

Computing power has increased exponentially since the law was first proposed in 1965. For instance, current microprocessors are almost four million times more powerful than the first microchip processors introduced in the early 1970s (Schatsky et al., 2014).

Nevertheless, there are indications we may be reaching the technological limits of Moore’s Law. Raw computing power (as measured by transistors per chip) is reaching something of an inflection, leading many to speculate we are approaching the “limits of Moore’s Law” (Simonite, 2016; and *The Economist*, 2016a). The number of transistors

per chip has been plateauing since the early 2000's (Figure 2). By Intel's own estimates, the number of transistors on a microchip may only continue doubling over the next five years (Bourzac, 2016).




Figure 1: Number of Transistor Components per Chip

Source: Moore (1965)

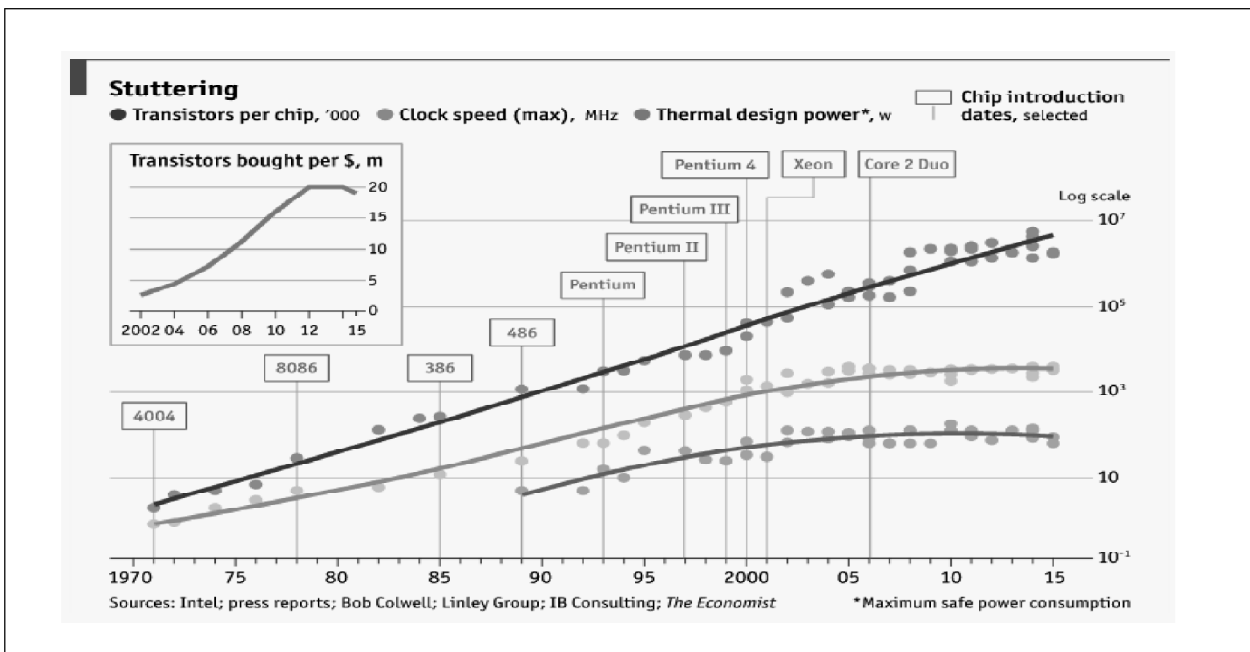


Figure 2: Computer Processing Speeds

Source: The Economist (2016a)

Chip manufacturers are approaching the theoretical limits of space and physics that makes pushing Moore's Law further both technologically challenging and cost prohibitive. Moore's Law became a self-fulfilling prophecy because Intel made it so. They pushed investment and catalyzed innovation to produce more power and faster

processing (*The Economist*, 2016). In the face of increasingly high costs and complex design considerations, processing speeds are unlikely to continue to grow in the same fashion.

While important, Moore's Law represents only one of several assessments of computing power. Other industry measurements capture different aspects of raw hardware power. One measurement, Floating Point Operations per Second (FLOPS), is a raw estimate of the number of calculations a computer performs per second, an indication of computational performance. Another, Instructions Per Second (IPS), estimates how rapidly computers can respond to specific instructions and inputs, providing an indication of processing speed.

The literature has attempted to estimate (in rough terms) global computing capacity using IPS and FLOPS as standard measurements. Hilbert and Lopez (2012) using a variety of data from 1986 and 2007, estimated global computing capacity to be around 2×10^{20} IPS. They also estimate growth rates for general purpose computing hardware to have been around 61% over the same timeline. In another longitudinal study, Nordhaus (2001) calculated that computing performance has improved at an average rate of 55% annually since 1940, with variation by decade. A study from Oxford University in 2008 estimated that since 1940, MIPS/\$ has grown by a factor of ten roughly every 5.6 years, while FLOPS/\$ has grown by a factor of ten close to every eight years (Sandberg and Bostrom, 2008).

Building on this literature, in 2015, contributors to AI Impacts, an open-source research project based at the Oxford Futures Institute, estimated global computing capacity to be something in the region of $2 \times 10^{20} - 1.5 \times 10^{21}$ FLOPS. But how does this power compare with the human brain? Plausible estimates of human brain computing power ranged from 10^{18} , 10^{22} , and 10^{25} FLOPS (Sandberg and Bostrom, 2008; and AI Impacts, 2016). In his 2005 book, Google's Ray Kurzweil claimed the human brain operated at the level of 10^{16} FLOPS. By these estimates, global hardware processing power has surpassed the human brain. Already, some of the most powerful supercomputers can process data in greater volumes and with much more speed than the human brain. Yet the human brain remains vastly more efficient, requiring only enough energy to power a dim light bulb, while the energy required for the best supercomputers could power 10,000 light bulbs (Fischetti, 2011).

3.2. Software Capabilities

AI development is being catalyzed by more than just more powerful hardware. Improved software has facilitated AI development is being catalyzed by more than just more powerful hardware. Improved software has facilitated the development of more complex and powerful algorithms, an essential component of many new AI technologies. Deep learning, software capable of mimicking the brain's neural network, can learn and train itself to detect patterns through exposure to data (Hof, 2013). Deep Learning technologies diverge from classic approaches to AI, which typically relied on a pre-programmed set of rules defining what machines "can" and "cannot" do. Deep learning is not constrained by established rules and has the capacity to "learn", but it requires vast amounts of data and often breaks down if there are frequent shifts in data patterns (Hawkins and Dubinsky, 2016). As shown in Figure 3, revenues from software using deep learning technology could reach over \$10 bn by the mid 2020's, up from just over \$100 mn in 2015 (Tractica, 2016). Deep Learning technology has enjoyed a renaissance alongside the growth of "Big Data," powered by the accessibility and penetration of the internet, mobile devices, and social media, among other things. The vast amount of data being produced in these areas helps improve the quality of machine learning algorithms, which can be "trained" through exposure to varied datasets (Guszcza et al., 2014).

While deep learning places a premium on data mining and pattern recognition, another emerging approach, Reinforcement Learning, moves toward decision-making and away from pattern recognition (Knight, 2017). Under this approach, AI machines "learn by doing"; that is they attempt to perform a specific task hundreds or even thousands of times. The majority of attempts result in failure, yet with each success, the machine slowly learns to favor behavior accompanying each successful attempt. Reinforcement Learning builds on behavioral principles outlined by psychologist Edward Thorndike in the early 20th century. He designed an experiment that placed rats in enclosed boxes from which the only escape was by stepping on a lever that opened the box. Initially, the rats would only step on the lever by chance, but after repeated trials they began to associate the lever with an escape from the box, and the time spent in the box fell sharply (Knight, 2017). In March 2016 AlphaGo, a Google program trained through Reinforcement Learning, defeated Lee Sedol, one of the world's top Go players. This result was especially surprising because Go is an extremely complex game that cannot be reproduced by machines with conventional or simple rules-based programming. Experts had thought that a machine wouldn't be able to defeat a human Go player for another decade or so (Knight, 2017).

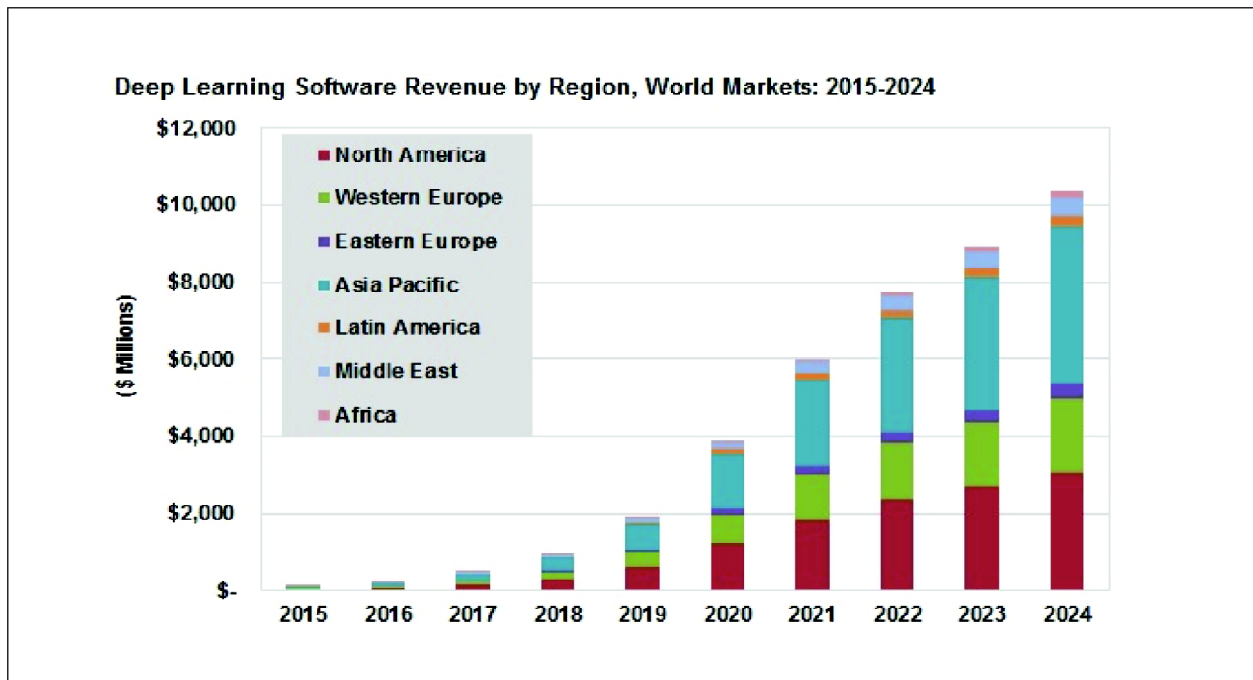


Figure 3: Forecasted Revenue for Software Built Using Deep Learning

Source: Tractica (2016)

3.3. Cloud Computing

Alongside Big Data, the internet and cloud computing (internet-based computing services) are important catalysts of AI development. They have helped to make vast amounts of data available to any device connected to the internet and they allow for crowdsourcing and collaboration that can improve AI systems (Schatsky et al., 2014). Cloud computing is fundamentally restructuring the licensing and delivery of software, operating platforms, and IT infrastructure. As shown in Table 1, it is catalyzing a movement towards providing software resources as on-demand services (Diamandi et al., 2011).

Cloud computing is still largely in its nascent stages, but the technology is evolving in parallel with many narrow AI applications. Microsoft’s website now offers many cognitive services through the cloud, including computer vision and language comprehension. Amazon Web Services has added data mining and predictive

Table 1: Cloud Computing Services

Computing Service	Description	Example Products
Infrastructure as a Service (IaaS)	Provides computing capabilities, storage, and network infrastructure.	Amazon EC2 and S3 Services Xdrive
Platform as a Service (PaaS)	Provide platforms that enable application design, development, and delivery to customers.	Microsoft Windows Azure Salesforce.com platform
Software as a Service (SaaS)	Software applications are delivered directly to customers and end users.	Google Docs Microsoft Office 365 Zoho

Source: Diamandi et al. (2011)

analytics tools as part of its cloud computing toolkit (Amazon, 2017). In 2015, telecommunications company Cisco released a white paper on the size and trajectory of global cloud computing capacity between 2015 and 2020. According to their estimates, global cloud IP traffic will grow at a Compound Annual Growth Rate (CAGR) of 30 percent between 2015 and 2020 (Cisco, 2016). They forecast annual global cloud traffic to reach 14.1 zettabytes (ZB) (1.2 ZB per month), by 2020, up from 3.9 ZB in 2015.²

Market spending on cloud computing services is projected to reach more than \$200 bn by 2020, up from an estimated \$122 bn in 2017 (IDC, 2016). Approximately 90% of global enterprises will use some type of cloud-based technology by 2020 (EIU, 2016). Despite the forecasted growth, a 2016 study from the Economist Intelligence Unit found that cloud computing, measured by industry adoption rates, is really only just beginning. The study surveyed leaders from five major industries (banking, retail, manufacturing, healthcare, education), and found that an average of only 7% of respondents felt that cloud computing played a “pervasive role” (Economist Intelligence Unit, 2016, p. 3). In addition to varied rates of adoption, concerns over privacy, security, and flexibility remain. Companies deciding to adopt one cloud platform may find it costly or difficult to transfer their information to another provider (Economist, 2015). Improved regulation that allows companies and consumers to move data between different providers may enhance adoption rates. The growth of the cloud, both in terms of data management and market size is undeniable, but important challenges remain.

3.4. The Shifting Investment Landscape

AI advancement has traditionally been the product of universities and corporate research and development labs (e.g., IBM). Over the last few years, Silicon valley has moved major investments into AI. There is a growing appreciation and recognition of the social benefits and commercial value of narrow AI technologies, prompting interest from silicon valley and private start-ups. Major technology companies including Facebook, Google, and Microsoft have hired some of the best minds in AI and invested heavily (Albergotti, 2014; and Regalado, 2014). One reason technology companies have been able to attract the top talent away from research universities is that, in addition to comfortable compensation packages, these companies are sitting on vast amounts of user generated data that are increasingly essential to AI development. These data are not publicly available.

Private investment in AI has grown commensurate with the results and attention. One market research firm estimated private funding for AI (excluding robotics) to have grown from \$589 mn in 2012 to over \$5 bn in 2016 (CB Insights, 2017). There may be as many as 2,600 different companies operating in the AI sector as of 2016, with over 170 having taken off in Silicon Valley since 2014 (Byrnes, 2016). The robotics market alone could be worth close to \$135 bn by 2019 (Waters and Bradshaw, 2016).

3.5. Information and Communication Technology Access

Information and communication technology access is another important indicator of AI. ICT penetration rates, particularly mobile broadband, serve as an important baseline to justify investment into AI and give some indication of the technological depth of a society. Many AI applications over the near-term will rely on smart phones as a service delivery mechanism. The number of smart phones in the world was expected to grow, reaching over 6 billion by 2020 with much of the growth coming from the developing world, with an estimated 3.2 billion in use (Ericsson, 2016). The 2016 annual report by the International Telecommunications Union (ITU) provided a snapshot of global ICT connectivity:

- Globally, 95% of the population lives in an area covered by a cellular network; 84% of the population lives in an area with a mobile broadband network (3G or above), but only 67% of the global rural population has access to mobile broadband regularly.
- An estimated 3.9 billion people are not using the internet regularly, roughly 53% of the total. Internet penetration rates in developed countries are around 81%, while in the developing world they average approximately 41%, but only 15% in the least developed countries.
- An estimated one billion households have internet access: 230 million in China, 60 million in India, and 20 million across the 48 least developed countries.

² Zetabyte is equal to 10^{21} bytes. A byte is a unit of digital information, traditionally consisting of 8 bits. 8 bits represents the number of bits required to encode and save a single character of text in a computer.

As we can see from the statistics above, much of the developed world is covered by internet access and mobile broadband, but a general lack of access constrains many developing countries.

Taken together, the preceding list comprises important proximate drivers of AI development. In addition, the spread of AI technologies for commercial and personal use will be contingent on policymaking and industry adoption. Transparent policymaking is necessary to define the rules of AI and its use, but also to justify its adoption and continued investment in these technologies. How rapidly the private sector can integrate emerging AI technologies into its work cycle will further hinder or hamper adoption. With these trends and important drivers in mind, we shift to thinking about “intelligence” and how we might evaluate or assess generally intelligent machines.

4. Measuring and Evaluating AI

There is minimal doubt that AI is a “successful” field; new technologies and applications are emerging regularly (Hernandez-Orallo, 2017, p. 117). Almost all recent progress has been restricted to narrow AI sectors; the development of AGI machines remains a distant goal rather than an imminent reality. Scientists and developers in the field remain confident that AGI will be developed, though there is significant uncertainty as to the timeline.

Evaluating AI requires some basic consensus around standard benchmarks of progress and an understanding of what qualifies as AGI, at least from a definitional perspective. As we will see, there exists a great many definitions of “intelligence,” a growing number of tests and evaluation techniques used to assess machine intelligence, and some dispute around how we can (or should) accurately measure AGI.

Early researchers of AI were focused on developing generally applicable machines, that is those capable of solving a variety of problems otherwise requiring “intelligence” (Newell et al., 1959). Some researchers tried to design programs that would be capable of solving questions commonly found on human IQ tests, such as, the ANALOGY program which sought to answer geometric-analogy questions frequently found on intelligence tests (Evans, 1964). Ultimately however, the creation of generally intelligent machines was far more difficult than many predicted, leading to a stagnation in AI research in the 1960s and the 1970s. The pace of research also slowed as a result of what has become known as the “AI effect,” or the idea that as soon as AI successfully solves a problem, the technology is reduced to its basic elements by critics and thus is no longer considered intelligent (McCorduck, 2004). For instance, when Deep Blue beat chess champion Gary Kasparov in 1997, critics claimed that the machine resorted to brute force tactics, which were simply a function of computing power rather than a true demonstration of intelligence (McCorduck, 2004, p. 33). The result of the “AI effect” is that the standards for what constitutes true machine intelligence keep moving. These difficulties helped in part to shift the field toward the development of narrow technologies capable of achieving measurable and practical results (Hernández-Orallo, 2017, p. 120).

4.1. Evaluating Narrow AI

The growth of narrow AI technology means that most AI is now assessed according to a “task-oriented evaluation,” (Hernández-Orallo, 2017, p. 135) that is, according to its relative performance along task-specific, measurable outcomes. Today, all of the benchmarks along narrow the AI categories discussed below measure performance according to the completion of a specific task:

- The ability to translate text from one language to the other; or
- Identify a cat from a series of photos; or
- Accurately respond to specific questions from a human user.

Progress along these evaluations shows that AI is becoming more useful but doesn’t necessary suggest that AI is becoming more intelligent. Measuring and evaluating AI requires some classification and understanding of the major technologies that are shaping the field. The AI field is diverse and rapidly expanding and resists simple classification. Pulling together various threads from a wide range of research, we have identified six “categories” (Table 2) of AI technology generating new breakthroughs: computer vision, machine learning, natural language processing, robotics, the “Internet of Things,” and reasoning/decision-making. These six include both foundational AI technologies as well as important technologies emanating from them. While items on this list are neither exhaustive nor exclusive (see Box 1), they provide a framework to begin building the representation of AI in IFs.

Table 2: Technologies Comprising the Narrow AI Representation in IEs

Type	Description	Example Products
Computer Vision	Ability of computers to identify objects, scenes, activities in images.	Medical imaging, facial recognition, retail, and sales.
Machine Learning	Ability of computers to improve their performance through exposure to data without pre-programmed instructions.	Any activity that generates substantial data. Examples include: fraud detection, inventory management, healthcare, oil and gas.
Natural Language Processing	Ability of computers to manipulate, write and process language, as well as interact with humans through language.	Analyzing customer feedback, automating writing of repetitive information, identifying spam, information extraction and summarization.
Robotics	The branch of technology specializing in the design and construction of robots.	Unmanned aerial vehicles, cobots, consumer products and toys, select services, manufacturing.
Internet of Things/ Optimization	Networking of physical objects through of embedded sensors, actuators, and other devices that can collect or transmit information about the objects. Requires collecting data, networking that data, and then acting on the information.	Two main applications: the use anomaly detection and optimization. Specific applications in energy supply and demand, insurance industry and optimization of premiums, healthcare, public sector management.
Reasoning, Planning, Decision-making	This represents an area of AI research concerned with developing ability of machines to reason, plan, and develop decision-making capacity. We represent it as a general “spillover category” of machine reasoning, an essential element of AGI.	Limited modern applications and development. Some basic reasoning technology has been used to assist in proving mathematical theorems

Box 1: Areas of AI Research

There are many sub-disciplines and areas of study within the field of AI, many more than could be effectively captured in any modeling effort. The 2016 Association for Artificial Intelligence annual conference alone accepted submissions to over 30 different AI subfields. The six main categories of technology we have represented within narrow AI cover both foundational AI technologies (computer vision, machine learning, natural language processing, reasoning), as well as important technologies that are emanating from the field (robotics, internet of things). These areas are currently receiving significant attention, deep financial investment, and/or are fundamental for advancing the spectrum towards AGI.

We recognize these categories are neither exclusive nor exhaustive. To outline the diversity of research and development currently happening within the field, Table 3 below depicts other important areas of AI

Box 1 (Cont.)
technological development. Included in this list are the main disciplines according to <i>AI Journal</i> , one of the leading publications in the field (Hernandez-Orallo, 2017, p. 148).

Table 3: Other Major Areas of AI Research not Explicitly Captured by the IFs Narrow AI Representation

AI Subfield	Definition
Crowdsourcing and Human Computation	Algorithms that allow autonomous systems to work collaboratively with other systems and humans.
Algorithmic Game Theory	Research focused around the economic and social computing dimensions of AI.
Neuromorphic Computing	Mimic biological neural networks to improve hardware efficiency and robustness of computing systems.
Automated (Deductive) Reasoning	Area of computer science dedicated to understanding different aspects of reasoning to produce computers that are capable of reasoning completely.
Constraint Processing	Refers to the process of finding solutions amidst a set of constraints that impose conditions that certain variables must satisfy.
Knowledge Representation	Representing real world information in forms that a computer system can use to solve complex tasks.
Multi-agent Systems	Computer system composed of multiple, interacting, intelligent agents within one environment.
Planning and Theories of Action	Developing machines capable of “understanding what to do next” in the context of unpredictable and dynamic environments, often in real-time.
Commonsense Reasoning	Simulating human ability to make presumptions, inferences, and understanding about ordinary situations that they encounter on a day-to-day basis.
Reasoning Under Uncertainty	Concerned with the development of systems capable of reasoning under uncertainty; Estimate uncertain representations of the world in ways machines can “learn from.”

4.2. Benchmarking Progress in Narrow AI

In this section, we outline recent progress along the categories of narrow technology outlined above. Given the lack of standardized data on AI technology and development across time, these benchmarks are pulled from a variety of sources, including (but not limited to), media reports, market research estimates, government analyses, journal articles, and other independent analyses of the field. Table 4 provides a summary of progress along the identified categories of narrow AI technology and an initial AI index score (from 1-10) for each estimated by the authors. A justification for the initial score is elaborated in text below the table.

4.3. Thinking About Measuring AGI

There are many, varying, conceptual measurements for AGI. One example is the “coffee test,” under which a machine should be able to enter an ordinary and unfamiliar human home, find the kitchen, and make a cup of coffee ([Moon, 2007](#)). Along these lines, others have proposed that a generally intelligent machine should be able to

Table 4: Benchmarking Progress in Narrow AI Technologies		
Technology	Performance Benchmarks	2015 Index Score
Machine Learning	<p>1997: IBM Deep Blue defeats Gary Kasparov, a Grandmaster, in a game of chess.</p> <p>2011: IBM Watson defeats Jeopardy! champion. In the lead up to the contest, between December 2007 and January 2010, the precision of Watson's responses more than doubled. Precision measures the percentage of questions the system gets right relative to those it chooses to answer. In December of 2007, Watson answered 100% of Jeopardy! style questions with only 30% accuracy. By May of 2008, the accuracy of response improved to 46%, and by August of 2008 it was close to 53%. A year later in October of 2009 accuracy (with 100% of questions answered) hovered around 67%, twice the level in 2007 (Ferucci et al., 2010).</p> <p>2008-2012: NIST Machine Translation Scores. Chinese to English translation accuracy (as compared with a human translation) improved 28-34% between 2008-2012. Arabic to English accuracy scores improved from 41% to 45%. Less widely spoken languages scored less well: Dari to English 13% (2012), Farsi to English 19% (2012), Korean to English 13.6% (2012) (NIST, 2012).</p> <p>2013: First AI software passes the Captcha test (Metz, 2013). Captcha is a commonly used authentication test designed to distinguish humans and computers. Captcha is considered broken if a computer is able to solve it one percent of the time; this AI software solved it 90% of the time.</p>	3
Computer Vision	<ul style="list-style-type: none"> • 2010-2015: Stanford AI ImageNet competition. Image classification improved by a factor of 4 over 5 years. Error rates fell from 28.2% to 6.7% over that time. • In the same competition, object localization error rates fell from 45% in 2011 to 11% in 2015 (Russakovsky et al., 2015). • 2012: Google releases the "Cat Paper," about a machine capable of learning from unlabeled data to correctly identify photos containing a cat (Le et al., 2012). • 2014: Facebook's "DeepFace" team publishes results that claim its facial recognition software recognizes faces with 97% accuracy (Taigman et al., 2014). • 2015: Microsoft image recognition algorithms published an error rate of 4.94%, surpassing the human error threshold of 5.1% and down from error rates of 20-30% in the early 2000's. 	3
Natural Language Processing	<ul style="list-style-type: none"> • 2012-2014: Siri's ability to answer questions correctly improved from an estimated 75% to 82%. Over the same time period, Google Now's response accuracy improved from 61% to 84%. Siri's ability to interpret a question when heard correctly improved from 88% to 96%. Google Now similarly improved from 81% to 93% (Hughes, 2014). 	2

Table 4 (Cont.)		
Technology	Performance Benchmarks	2015 Index Score
	<ul style="list-style-type: none"> • 2015: Baidu released its DeepSpeech 2 program that can recognize English and Mandarin better than humans and achieves a character error rate of 5.81%. Represents a reduction in error rates by 43% relative to the first generation of the software. • 2016: Microsoft switchboard word transcription error rates have dropped from between 20-30% around 2000, to a reported 5.9% in 2016 (Xiong et al., 2016). 	3
Robotics	<ul style="list-style-type: none"> • 1942: Isaac Asimov publishes the Three Laws of Robotics. • 1954: Patent for “Unimate,” the first industrial robot filed. Unimate worked on a General Motors assembly beginning in 1961. • 1969: Robot vision for mobile robot guidance first demonstrated at Stanford. Hitachi develops the first robot capable of assembling objects from assembly plan drawings (International Federation of Robotics, 2017). • 1970: Hitachi develops the first robot capable of assembling objects from assembly plan drawings. • 1980: First use of machine vision in robotics demonstrated at the University of Rhode Island in the US. • 1990: Manufacturers begin to implement network capabilities among robots. • 2002: Reis Robotics patents technology permitting among the first direct interactions between humans and robots. The size of the robotics industry crosses \$1 bn. • 2003: Mars Rover first deployed heading to the planet Mars. Mars Rover missions continue through the present day. • 2004: First DARPA Grand Challenge. Goal: design an autonomous car capable of completing 150-mile route through the Mojave Desert in the US. No cars completed the route; an entry from Carnegie Mellon went the farthest, completing roughly 7.3 miles. • 2005: Second DARPA Grand challenge. Design a driverless car capable of completing a 132-mile off-road course in California. Of the 23 finalists, 5 vehicles successfully completed the course, the fastest in just under seven hours. • 2007: Third DARPA Grand Challenge. Design a self-driving car capable of completing an urban, 60-mile course in less than six hours. Required vehicles that could obey traffic laws and make decisions in real time. Six teams successfully completed the course, the fastest in just over four hours (DARPA, 2017). • 2015: Carmaker Tesla releases its first-generation Autopilot technology, part of its suite of self-driving technology. Autopilot 	1

Table 4 (Cont.)		
Technology	Performance Benchmarks	2015 Index Score
	<p>allows Tesla to automatically steer within lanes, change lanes, manage speed, and parallel park on command.</p> <ul style="list-style-type: none"> • 2015: The University of Michigan opens MCity, a testing center for autonomous vehicles. Represents the first major collaboration between private industry, government, and academia on the development of autonomous vehicles (Michigan News, 2015). • 2015: BCG estimates global robotics manufacturing installations to grow 10% through 2025, reaching an estimated 5 million globally. Yet even by 2025, robotics may only account for 25% of all manufacturing tasks globally (Sirkin et al., 2015). 	3
IoT	<ul style="list-style-type: none"> • 1990: There are an estimated 100,000 internet hosts across the worldwide web. • 2000: More than 200 million devices connected to the IoT • 2012: A botnet known as “Carnabot” performed an internet census and counted approximately 1.3 billion devices connected to the worldwide web. • 2014: The number of devices communicating with one another surpassed the number of people communicating with one another. • 2015: over 1.4 billion smart phones were shipped and by 2020 it was estimated there would be 6.1 billion smartphone users. • 2020: There could be anywhere from 20-50 billion devices connected to the IoT. 	2
Reasoning, Planning, and Decision-making	<ul style="list-style-type: none"> • Spillover category designed to capture progress towards reasoning, planning, and decision-making, key elements of general intelligence. • There are very minimal current applications in this technology. Automated Reasoning, for instance, has been used in the formal verification of mathematical proofs and the formalization of mathematics. 	1

enroll, take classes, and obtain a degree like many other college students ([Goertzel, 2012](#)). Nils Nilsson, a Professor of AI at Stanford, has taken the definition a step further, proposing an “employment test,” whereby a truly intelligent machine should be able to complete almost all of the ordinary tasks humans regularly complete at their place of employment ([Muehlhauser, 2013](#)).

These definitions of AGI have similar underlying themes: they require that machines be able to respond to different tasks under varying conditions. These differing tests help us arrive at a working definition of general-purpose AI systems, proposed by Hernandez-Orallo ([2017, p. 146](#)):

AGI must do a range of tasks it has never seen and not prepared for beforehand.

Having defined AGI, we must now consider measurement techniques. The Turing Test, first proposed by English Mathematician Alan Turing in 1950 has evolved into a simple test of intelligence. The Turing Test measures the ability of machines to exhibit intelligent behavior indistinguishable from that of humans. If a machine can fool a human into thinking it is human, then that machine has passed the Turing Test. Some have identified it as “a simple test of intelligence” (French, 2000, p. 115), or a goal of AI (Ginsberg, 1993, p. 9). An example of the enduring appeal of the Turing Test, the Loebner Prize for Artificial Intelligence, offered \$100,000 to the chatterbot deemed to be most “human-like” according to a panel of judges. The prize was offered annually between 1991 and 2020.

Box 2: What Capabilities Might Machines Need for AGI?

Some researchers of AI have proposed a suite of tests for which to analyze general intelligence. Adams *et al.* (2012) identified “high-level competency areas” that machines would have to depict across a number of scenarios, including: video-game learning, preschool learning, reading comprehension, story comprehension, and the Wozniak test (walk into a home and make a cup of coffee) (synthesized from Hernandez-Orallo, 2017, p. 148).

Core competency areas as identified by Adams *et al.* (2012) and reproduced in Hernandez-Orallo (2017) are in the table below:

Perception	Memory
Attention	Social interaction
Planning	Motivation
Actuation	Reasoning
Communication	Learning
Emotion	Modelling self/other
Building/creation	Use of quantities

Source: Adams et al. (2012)

While such a complex suite of assessments to measure AGI may never be possible across all of the competencies identified in Table 5, comprehensive analysis of AGI could include some combination of these different assessments.

More recent research has argued against the Turing Test as a sufficient measure for AGI. Hernandez-Orallo (2017, pp. 129-130), summarizes its shortcomings succinctly. He points out that many non-intelligent machines can be trained and designed to fool judges, without necessarily exhibiting true intelligence. The results of the Turing can differ dramatically based on indications, protocols, personalities, and intelligence of the people involved, including both the judges and the participants. Finally, the Turing Test asks machines to imitate humans, which raises questions about how representative the imitation is of the entire human race.

Instead of focusing on task-specific evaluations, AGI evaluation should move towards “feature-oriented evaluation.” Such an evaluation would be based on a profile of behavioral features and personality traits of the machine, rather than its ability to perform a discrete task (Hernandez-Orallo, 2017, p. 146). This type of evaluation builds on performance along narrow task areas and towards a maximalist view of AGI. The type and style of this evaluation is debated and ill-defined. Some have proposed the idea of a machine cognitive decathlon (Hernández-Orallo, 2017; and Vere, 1992), or a test of mental flexibility. Feature-oriented evaluation is complicated by non-specific questions around defining and measuring “personality.” Feature-oriented evaluation remains a nascent idea and topic, combining both measurements and evaluations of cognitive ability and personality (Hernandez-Orallo, 2017, p. 150), but it may be the direction the field moves toward with respect to assessments of AGI.

5. IFs: Representing AI

We now turn to a discussion of the construction and conceptualization of the AI indices in IFs. Understanding the IFs platform is important for understanding how the AI representation is integrated within the tool and how it

could be used to model the impact of AI. IFs is an open-source, quantitative modeling tool for thinking about long-term futures. Building on 5,000 historical data series, IFs helps users understand historical patterns, explore the current path of development and the trajectory we appear to be on (or the current path), and shape thinking about long-term futures. To do this, IFs leverages relationships across hundreds of variables from 12 dynamic, interconnected systems of human development. Figure 4 depicts the major sub-modules of the IFs system. The linkages shown are illustrative rather than comprehensive; each link is comprised of hundreds of variables. The IFs current path represents expectations of how development will unfold across each of these systems absent significant alteration or intervention (think drastic policy change, man-made or natural disasters, conflict, or technological discontinuities). The current path provides a necessary reference point for alternative scenario analysis. It is itself a dynamic forecast, driven by the variables and relationships built into the model. Many of the assumptions in the model can be modified by users to better reflect their own understanding of how these systems are developing and unfolding across time.

IFs is developed by the Frederick S Pardee Center for International Futures, based at the Josef Korbel School of International Studies at the University of Denver in Colorado, USA. It is available at <https://korbel.du.edu/pardee> for both on-line use and download. See that website and Hughes (2019) for extended information on IFs.

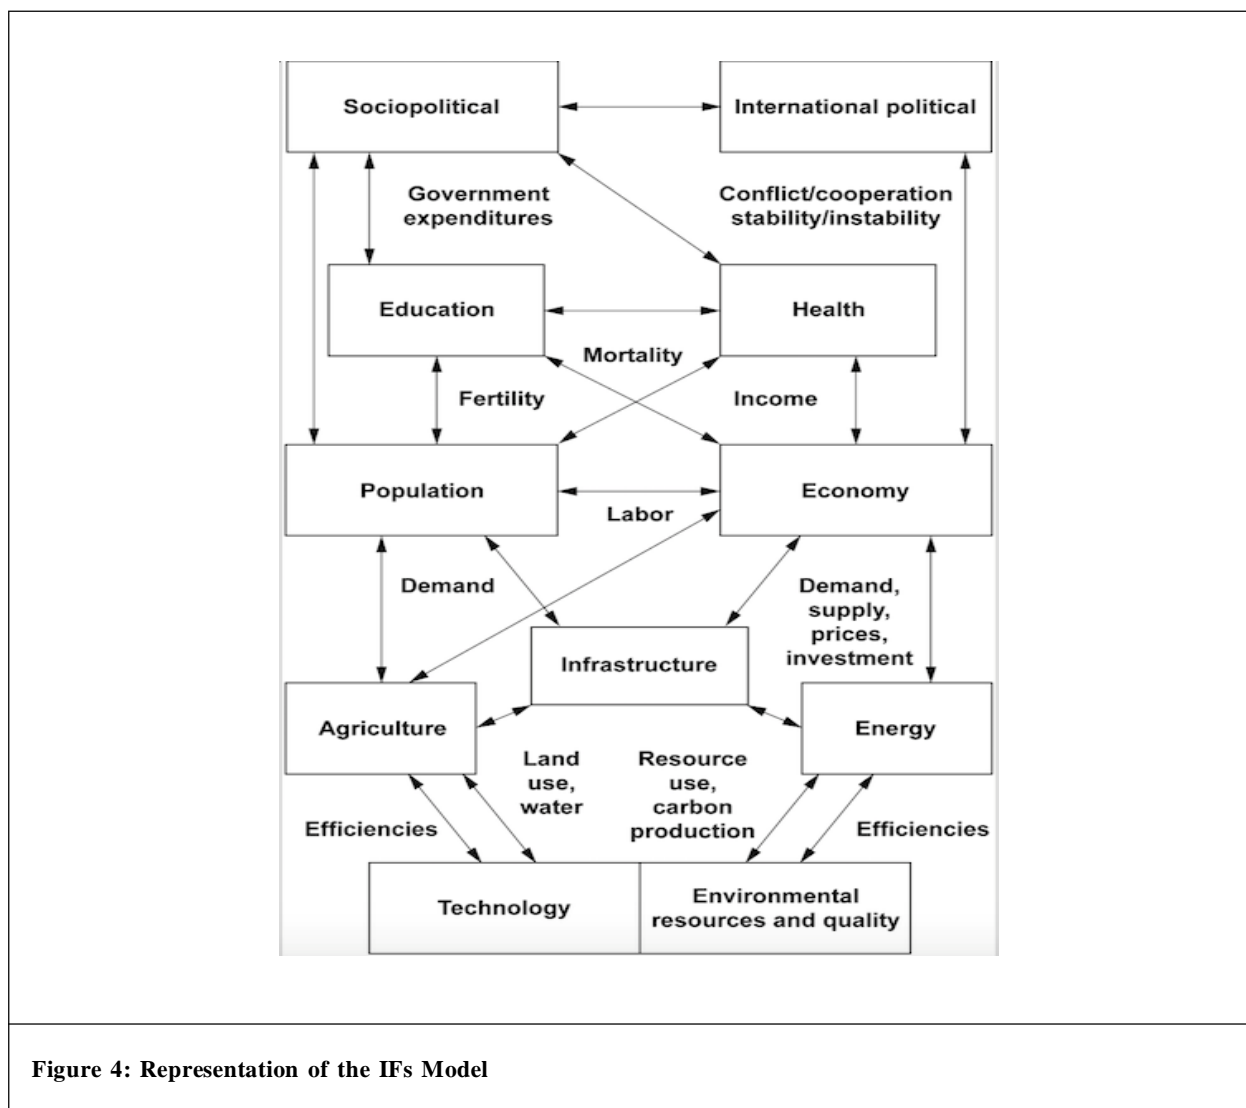


Figure 4: Representation of the IFs Model

5.1. AI Variables in IFs

The AI forecasting capability in IFs is a set of indices that estimates and forecasts global development of AI. At present it does not contain forward linkages, a task we discuss in later sections of this paper. We have added several variables to the IFs platform to develop the modeling capability. The AI representation forecasts progress along narrow, general, and super AI consistent with the conceptualization discussed earlier.

The first variable added to IFs, AITASK, estimates and forecasts technological progress along each of the six areas of narrow AI we defined above: computer vision, machine learning, natural language processing, Internet of Things, robotics, and reasoning. AITASK is represented as an index scaled from 1 to 10, where 1 represents no development, and 10 represents full or complete development (see below for a more in-depth discussion of our thinking in this respect). The index score along each of these narrow technologies is initialized in 2015 (the IFs base year).

The second variable added to IFs, AITASKGR, represents the annual growth rate along each of these technologies, and saturates on approach to 10 for each. Each narrow technology grows at a different pace, estimated by the authors using inputs like performance benchmarks, complexity of each technology, investment, and levels of research. AITASK Reasoning grows at the slowest pace of the AITASK indices. Progress along this index represents the movement towards machines capable of reasoning completely, complex decision-making, and provided with a sense of purpose and awareness of the world around them. Any movement from narrow to AGI in the IFs index is implicitly constrained by the pace of AITASK Reasoning, regardless of progress among the other areas of narrow AI development.

Finally, we have also added AIMACHIQ, a variable which represents the movement from narrow AI to general and superintelligent AI. AIMACHIQ is scaled as an index representing machine IQ scores, roughly corresponding with human-level IQ scores. In the current path, the movement from narrow to AGI occurs when an index score of 10 is achieved for each of the narrow technologies denominated under AITASK, except for AITASK Reasoning, which is at 5. This transition is reflected on AIMACHIQ at an index score of around 60. At that point, the index forecasts AGI will have been achieved, though a score of 60 corresponds to machines with the equivalent of low-level human intelligence. AIMACHIQ then grows algorithmically as AITASK Reasoning continues to improve, saturating toward an index score of 200 as AITASK Reasoning reaches 10. An AIMACHIQ score of between 180 and 200 represents machine superintelligence, as this would correspond with some of the highest reported IQ scores among humans.³

Table 6: AI Variables Added to IFs to Operationalize the AI Forecast

		Definition	Scale
Variables	<i>AITASK</i>	Index measuring developmental progress of six areas of narrow AI technology: machine learning, computer vision, natural language processing, IoT, robotics, and machine reasoning. IFs forecasts development along each of these different narrow technologies	1-10 (for each category of narrow AI technology)
	<i>AITASKGR</i>	Represents estimated, differential, annual growth rates of each narrow technology.	1-200
	<i>AIMACHIQ</i>	Index measuring the level and capacity of machine intelligence. Index scores correspond approximately to human-level IQ scores and intelligence.	
Parameters	<i>aitaskm</i>	Multiplicative parameter allowing users to adjust the growth rate of task-specific technologies. Users can accelerate or slow this parameter by up to 1,000% in either direction.	Set to 1 in the Current Path
	<i>aimachiqm</i>	Multiplicative parameter allowing users to adjust the growth rate of general and superintelligent AI. Users can accelerate or slow this parameter by up to 1,000% in either direction.	Set to 1 in the Current Path

³ Marilyn Vos Savant has the highest living recorded IQ today with a score of 228. Renowned physicist Stephen Hawking has a recorded IQ of around 160.

In addition to each of the variables, we have added parameters described in Table 6 to each of the AI variables. Parameters allow users to exogenously adjust the AI representation with maximum flexibility to bring the forecast in line with their own expectations of AI development.

There is no comprehensive, standardized dataset or series of benchmarks measuring the growth of AI from which we can draw. There is also much debate and controversy over the pace of development and uncertainty around what the future of the field could look like. With that uncertainty in mind, the next section outlines the thinking behind the indices and growth rates along the six categories of narrow AI technology modeled within IFs.

5.2. Initializing AITASK: Rapid Progress over the Past Five Years

Many of the notable performance benchmarks outlined in Table 4 have occurred recently. If we were constructing this AI forecast 5 to 10 years ago each of these technologies would have been initialized with a score of one. New breakthroughs in Deep Learning technology, a foundational element of many of the technologies above, including computer vision, machine learning, and natural language processing, has been responsible for much of the progress. Deep Learning and artificial neural network technology has been around since the 1980s and 1990s but operated largely at the fringes of main AI research.

Today however, the results produced through Deep Learning have come about because researchers have the means to store, manipulate, and utilize the vast amount of data produced by an increasingly digital world. The result has been an explosion of successful technologies. Stanford's ImageNet competition began in 2010. Apple iPhone's automated assistant Siri was acquired in 2010 and first introduced as part of the iPhone product line in 2011, Google responded by releasing Google Now in 2012. Google Brain, the project at Google centered on Deep Learning, opened in 2012. According to a company spokesperson, in 2012 Google was working on two Deep Learning projects. By 2016 it was working on over 1,000 (Parloff, 2016). In 2016, Google overhauled Google Translate using artificial neural networks, showing significant results in both accuracy and fluency of translation. These improvements were the result of a project that began in 2011. In 2013, Facebook hired Yann LeCun, a leading Deep Learning scientist, to run its new AI lab. In 2016 Microsoft consolidated much of its AI portfolio into an umbrella AI and Research Group, which brings together more than 5,000 computer scientists working on AI-based projects (Microsoft, 2016). According to CB Insights, a market analytics firm, in the second quarter of 2016 nearly 121 rounds of equity fundraising were held for AI-based start-ups, compared with just 20 in 2011 (Parloff, 2016).

5.3. Initializing AITASK: Understanding the Shortcomings of Today's Technology

Yet, despite some references to the recent period as the "the Great AI Awakening," (Lewis-kraus, 2016), the functionality of AI remains very limited. As AI pioneer and Director of Baidu AI, Andrew Ng, points out, almost all AI technologies today operate on a simple premise: data input is used to generate a simple response (Ng, 2016). In this section we look at the current shortcomings of each AI technology to provide context for and justify the initial indices score.

5.4. Machine Learning

5.4.1. AITASK Machine Learning 2015 Index Score: 3

New algorithms that improve both the accuracy and speed of machine learning have been fueled by new advances like Deep Learning and Reinforcement Learning. Corresponding performance in task-specific activities reflects that improvement (reflected in Table 4). Additionally, the market for machine learning technology was estimated at around \$613 mn in 2015, forecast to grow to \$3.7 bn by 2021 (Markets and Markets, 2016a), suggesting these improvements are catalyzing interest and funding. Yet many improvements have not necessarily been uniform. For instance, machine translation accuracy is much lower among less commonly spoken languages. In 2012, the accuracy of Korean-to-English translation or Farsi-to-English translation hovered between only 13 and 19%, while it had improved to over 35% for Arabic and Chinese translations. Machine learning technology today remains dependent on massive volumes of data to "train" machines. Humans must be involved in the production, manipulation, and management of the data. Examples of common applications of machine learning are listed in Table 7. Each involves a simple binary output and massive data input. While each is a simple task for a human, as we will see below, machines can be easily fooled.

A result of these benchmarks, we have initialized AITASK Machine Learning at 3 in 2015. A machine learning index score of 10 represents perfect machine learning capabilities. To achieve an index score of 10, machine

Table 7: Examples of Machine Learning		
Input A	Output B	Application
Picture	Does the picture contain faces? (0,1)	Photo tagging
Loan application	Will the user repay the loan (0,1)	Finances
Add and user information	Will this user click on the ad? (0,1)	Ad-based targeting

learning would be capable of learning almost any task as well as a human, with the ability to produce complex, sophisticated output. Additionally, machine learning approaching 10 would contain sophisticated algorithms such that it is capable of learning from far smaller volumes of data than today's models. That technology might even be able to manipulate and absorb data without human input.

5.5. Computer Vision

5.5.1. AITASK Computer Vision 2015 Index Score: 3

Another area which has seen rapid improvement in the last five years is computer vision. The AI ImageNet competition, hosted by Stanford, has reported significant improvement in image identification, localization, and object detection between 2011 and 2015 (see Table 4). The market for computer vision is estimated to grow from \$5.7 bn in 2014 to over \$48 bn in 2022 (Tractica, 2016).

But it still remains fairly easy to fool computers into seeing something that is not there, or misclassifying objects erroneously. Many of the tasks recently completed by computer vision are extremely basic for humans. The remaining important differences between machine and human vision that scientists do not fully understand and thus cannot build in a machine. Machines can still be easily fooled in ways that human vision would not be. A 2015 paper found that it was fairly simple to produce images that humans would immediately identify as gibberish, only for a computer to classify them as objects with 99% confidence (Nguyen et al., 2015). Another similar study found that changing images in ways almost imperceptible to humans caused machines to misclassify objects entirely, in one instance classifying a lion as a library (Szegedy et al., 2013). More recently, researchers in France and Switzerland showed that small, almost imperceptible changes to an image could cause computers to mistake a squirrel for a fox, or a coffee pot for a macaw (Moosavi-Dezfooli et al., 2016; and Rutkin, 2017).

These challenges stem from fundamental differences in the way that humans and computers learn to “see” images. Children in school learning to recognize numbers eventually learn to recognize the common characteristics of each number after seeing many different examples. Ultimately, they come to recognize numbers even if the way the numbers are written is new to them. Computers learn to see by being fed millions of images of labeled data. The computer picks up the features that enable it to correctly identify the object of interest. But, machines, unlike humans, cannot see the whole picture. They learn from the pixels in a photo, while learning how tell different pixels apart. Therefore, minor changes in the pixel composition, alterations that do not change the image in the photo and would not fool a human, could fool a machine into thinking the photo is something that it is not (Rutkin, 2017).

Given the rapid progress in image and object identification, but accepting the existing significant limitations, we initialize AITASK Computer Vision at an index score of 3 in 2015. A computer vision index score of 10 would reflect computers with vision on par with humans, with the ability to distinguish, localize, differentiate without being easily fooled. Building machines with vision equivalent to that of a human also requires elements of reasoning to be able to identify, process, and understand the world they “see.”

5.6. Natural Language Processing

5.6.1. AITASK Natural Language Processing 2015 Index Score: 2

Natural language processing has improved both in terms of its ability to answer human-generated inquiries and its ability to decipher and translate between different human languages. Investment in and attention to the market have both increased; the market for natural language processing products is expected to grow from \$7.6 bn in 2016 to \$16 bn by 2021 (Markets and Markets, 2016b).

Arguably however, language remains one of the final frontiers of human intelligence. Machines capable of a full suite of natural language capabilities are still more of a distant dream than a short-term reality. Machines still do not “understand” language. Their ability to produce accurate, automated translation from spoken word in real time is limited by challenges that humans navigate with ease. Individual sounds are often not pronounced or spoken in isolation, in regular human conversation they are delivered in a constant stream. Machines still have difficulty in understanding nuanced vocabulary, children and elderly speakers, or competing with significant background noise (*The Economist*, 2017).

Researchers are also interested in producing machines capable of speech generation and conversation. The use of artificial neural network technology has helped researchers develop machines capable of producing more fluent sounding speech, but speech generation represents a whole new set of complex challenges. For instance, prosody, the modulation of speed, pitch, and volume to convey meaning, is an important component of human speech and interaction, which computers lack. Developing computers able to place stress on the correct words or parts of a sentence to convey meaning is incredibly difficult, and likely only “50% solved” by one estimate (*The Economist*, 2017). Additionally, fluent conversation is built around shared knowledge and an understanding of the world, something that machines lack. In theory, conversation between humans and machines represents a series of linked steps: speech recognition, synthesis, analysis of syntax and semantics, understanding of context, and dialogue, as well as common-sense and practical real-world understanding. Scientists still do not fully understand how the human brain pulls all of these disparate threads together to generate conversation; doing so in machines is a long-term task (*The Economist*, 2017).

Natural language processing is initialized at index score 2 in 2015. Fully automated machine transcription and translation remains a distant dream. Language is often considered the defining frontier of human intelligence. The Winograd Schema challenge, designed specifically to test how well machines understand and interpret language, was first held in 2016. The best entry scored a 58%, a result described as a “bit better than random” (Ackerman, 2016). According to some, machine transcription, translation, or language generation will never replace the benefits derived from understanding language and human-led translation. When people learn new words and phrases, they are not just learning the literal semantics or syntax of the individual words, they also learn cultural values and norms (Lewis-kraus, 2016).

A score of 10 along the natural language processing index represents machines capable of fully automated transcription and translation with close to 95% accuracy (roughly human level). A score of 10 represents machines capable of hearing, understanding, synthesizing, and generating language to participate in complex conversations on a variety of topics for which it has not necessarily been trained.

5.7. IoT

5.7.1. AITASK IoT 2015 Index Score: 2

The growth of the IoT has been fueled by rising internet connectivity and mobile technology penetration. Smart phones in particular are essential, as a service delivery and data collection mechanism and will remain one of the primary interfaces through which users interact with the IoT. The IoT has been and is forecast to continue growing exponentially, by some estimates there could be as many as 50 billion devices connected to the IoT by around 2020 (Howard, 2015; Figure 5).

Despite the sheer growth in the number of devices connected to the IoT, the technology is still very much in its infancy. The rules and norms that govern the use of and privacy around IoT-generated data remain ill-defined and opaque. Maximizing the benefits of IoT data requires interoperability between different IoT systems, today the vast majority of these systems are not interoperable. Finally, most data generated by the IoT today is used for basic tasks like anomaly detection and control, rather than for service optimization or predictive analytics, its most useful function (Manyika et al., 2015).

For these reasons, the IoT index is initialized at 2 in 2015, but is forecast to grow rapidly given the anticipated exponential growth in the number of connected devices. An index score of 10 represents a world where IoT data is protected and privacy concerns assuaged. Data produced is harnessed and analyzed to maximize efficiency on a broad social level. Fully smart cities and smart homes are the norm in most major developed urban areas. Automated transportation has become widespread not only as a result of autonomous vehicles, but also because cities are investing in the sensors and technology needed to produce the smart infrastructure that supports automated

driving. Smart infrastructure could include sensors embedded in the roadway that manage the flow and speed of traffic, sensors at intersections to reduce accidents and congestion, and smart lanes capable of charging cars as they drive (Manyika et al., 2013). According to a common definition of “smart” technology, global spending on smart city technology could cumulatively reach \$41 tn over the next 20 years (Pattani, 2016).

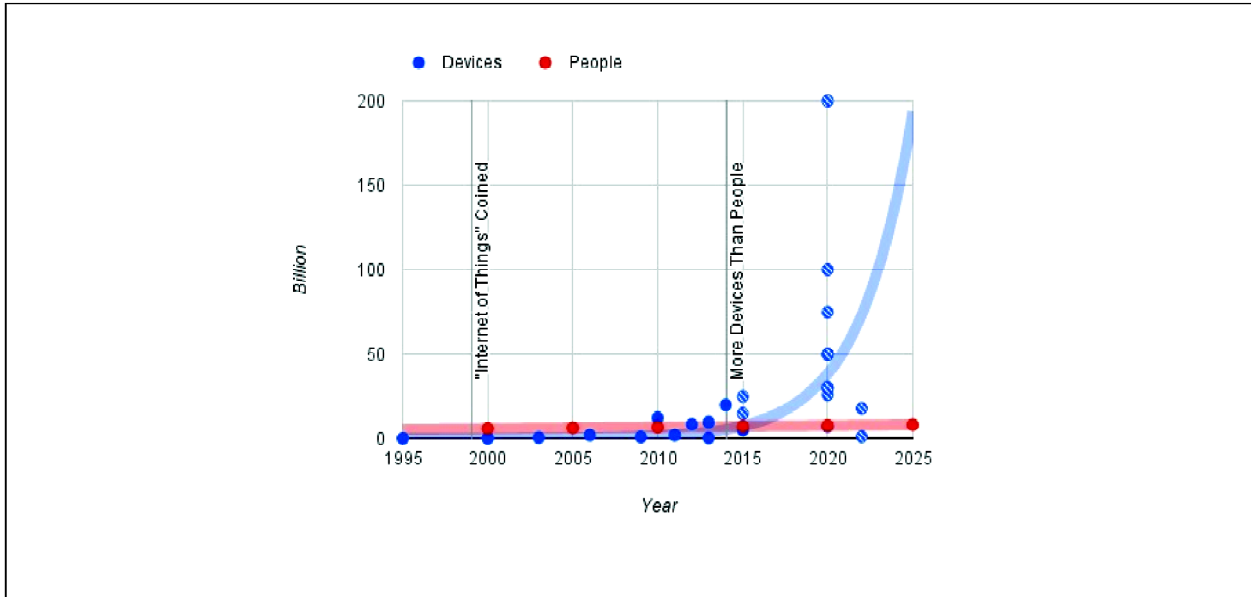


Figure 5: Number of Devices Connected to the Internet of Things vs. Size of the Population

Source: Howard, 2015

5.8. Robotics

5.8.1. AITASK Robotics 2015 Index Score: 1

Robots are already well-established in a number of fields, particularly manufacturing. According to a 2015 report by Boston Consulting Group, robots accomplish close to 10% of the tasks in the manufacturing industry today. Between 2010 and 2015, industrial robotics sales increased by a compound growth rate of around 16% annually, in 2015 there were 254,000 industrial robots sold (International Federation of Robotics, 2016).

The field of robotics is initialized at an index of 1 in 2015. This might seem surprising, given the large swaths of manufacturing and light industry jobs already replaced by robots (Frey et al., 2016; Frey and Osborne, 2013; and Schwab and Samans, 2016a). The functionality of most modern robots, however, remains limited. Robots today can perform a significant number of basic tasks that humans no longer want to do (particularly in manufacturing), or a few select tasks that humans cannot perform, (such as traversing the surface of Mars). The field is moving towards the creation of robots that are capable of working efficiently and effectively alongside humans. These so-called “cobots,” have proved difficult to make and account for roughly 5% of total global sales (Hollinger, 2016).

Robots cannot complete tasks they were not constructed specifically to undertake. In addition, robotics technology builds on other areas of narrow AI like computer vision, machine learning, and natural language processing. Robotics brings together both hardware and software, advancing the field of robotics requires improvements in both domains. Available market research suggests that investment is coming. One estimate placed the global robotics market at around \$71 bn in 2015, growing to \$135 bn by 2019 (Waters and Bradshaw, 2016). The size of the service robotics market alone could grow from around \$9 bn in 2016, growing to \$24 bn by 2024 (Zion Market Research, 2017).

An index score of 10 would be a robot that can respond to and perform a wide range of tasks for which it has not formally prepared or trained. A score of 10 may even represent a robot that can perform any general task as well as a human. This remains a distant goal. For instance, in 2016 Amazon held a contest to design a robot capable of stocking shelves in its warehouse. A task that would be fairly simple with humans, the winning robot had an error rate of around 16%, and Amazon said that they did not plan to make human workers redundant in spite of these results (Vincent, 2016).

5.9. Reasoning, Planning and Decision-making

5.9.1. AITASK Reasoning 2015 Index Score: 1

This variable is initialized at 1 in 2015. Development along this index is a distal driver pushing narrow AI technology toward the AGI level. Along this index, as reasoning approaches a score of 5, we forecast low-level and basic generally intelligent machines to begin to come into being. As the index moves towards 10, AGI improves, approaching the intelligence and capabilities of the average human. A reasoning score of 10 corresponds to the advent of a generally intelligent machine on par with human capabilities in reasoning, planning, language, vision, and decision-making. At this point, machine technology has a sense of purpose and understanding of the world around it.

6. Preliminary Results and Discussion

We begin by presenting the current path (or base case) results of the IFs AI forecast. Figure 6 shows the forecast of narrow AI technology along the six key technologies. The rate of development is calculated and estimated as a function of performance along task-specific competitions and evaluations, the estimated size of the market for each of these technologies and forecasted growth of that market, as well as (where available) estimates of academic publications in each of these domains. The IoT reaches an index score of 9 first, around 2038. Computer vision also proceeds rapidly, reaching an index score of between 9 and 10 around 2040. Robotics and natural language processing are slower-moving, and do not reach a score of 9 or 10 until around 2050.


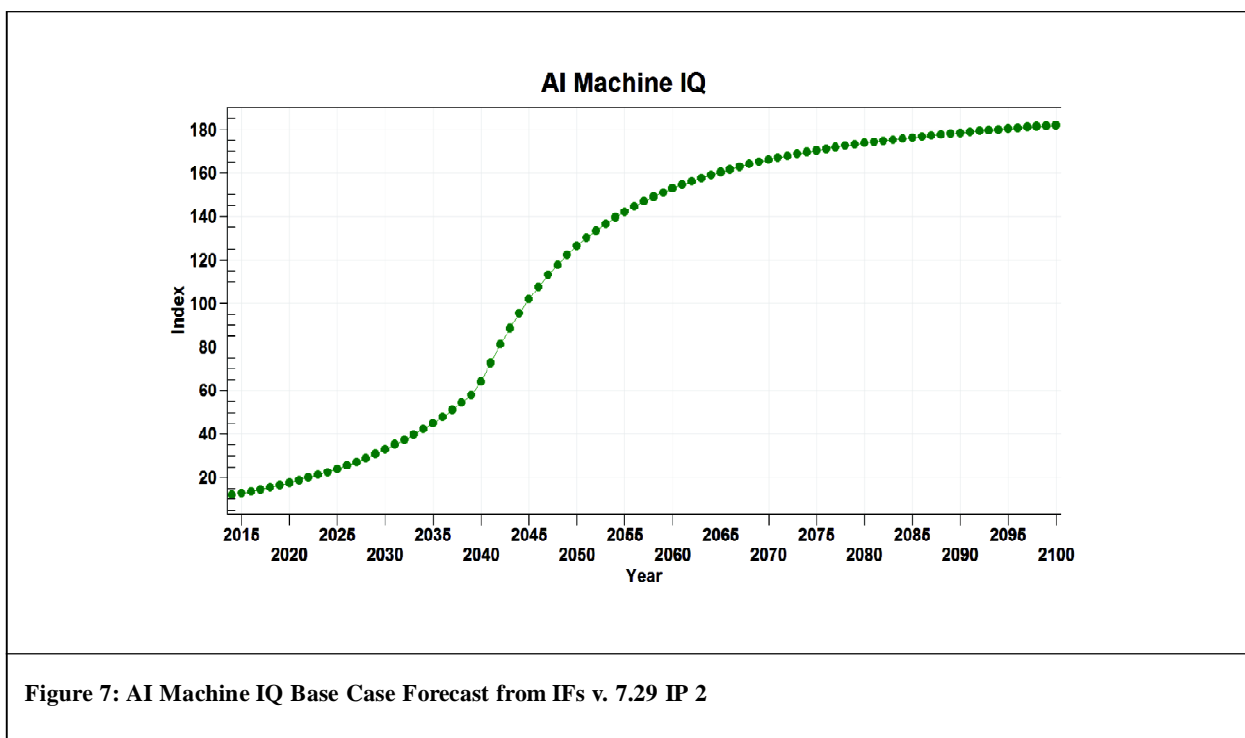


Figure 6: Narrow AI Forecast from IFs v. 7.29 IP 2

Under this approach, the movement from narrow AI to AGI is conceived of from a “bottom-up” perspective. The emergence of a generally intelligent machine must be developed from and build on existing narrow technologies. AGI researchers have expressed support for this approach (Harnad, 1990), and from our perspective this is conceivably the only way that AGI is likely to emerge. Progress along each of these technologies proceeds at differential rates, and AGI will not emerge until these technologies have reached advanced levels and become more integrated. Moreover, progression towards AGI is constrained by the movement of AITASK Reasoning, which is both the least developed and slowest moving of each of the narrow technologies. AGI is achieved when the reasoning index reaches a score of 5, which corresponds with a machine IQ score of between 55 and 60, or that of a human with very low intelligence. Figure 7 shows the Current Path forecast of AIMACHIQ. The current path suggests that a generally intelligent machine could be developed as early as 2040, although such a machine would have the intelligence equivalent to that of a “low-intelligence” human. AIMACHIQ suggests that a generally

intelligent machine with average level human intelligence (generally considered an IQ score between 90 and 110) could more likely be achieved between 2046 and 2050.

From there, AIMACHIQ is forecast to grow, in line with improvements in the capability of AGI. AIMACHIQ approaches a machine IQ score of 144, the equivalent of a high-intelligence score on the human IQ index by between 2055 and 2057. AIMACHIQ begins to approach super-human IQ (around 180, which only a handful of known humans have ever achieved) by 2090, suggesting that superintelligent AI could be achieved (at the earliest) near the end of the current century.



We fully acknowledge the vast amount of uncertainty surrounding the development of AI and the variability around a potential timeline. No comprehensive roadmap for AGI exists. The best available estimates of when we may see AGI come from expert surveys from the field. These provide important context for the IFs current path forecast.

The results from a number of studies using the Delphi Technique⁴ on the future of AGI are depicted in Table 8. The majority of respondents felt there is a 50% chance of AGI between 2040 and 2050, and a 90% chance of AGI on or after 2075. Notably, in one survey close to 2% of respondents felt that AGI would never be achieved.

In addition, Mueller and Bostrom (2014) also asked participants when they felt that we were likely to see the transition from AGI to artificial superintelligence. The responses indicated a 10% likelihood that the transition could occur within two years of the development of AGI and a 75% likelihood within 30 years of AGI. The IFs forecast is generally in line with these expectations.

We also created several scenarios around the future of AI development using the parameters described in Table 6. The scenarios are Accelerated AI and Stalled AI. Under the Accelerated AI scenario, AI development proceeds at roughly double its pace relative to the current path. In this scenario, AGI emerges around 2030, and superintelligent AI is forecast to emerge midway through the current century. Under the Stalled AI scenario, the reverse is true and AI development proceeds at half the pace of the Current Path. AGI is not forecast to emerge before approximately 2051, and superintelligent AI is not achieved within this century. Even by close to 2100, the available AI technology has a measured IQ score of around 90, on par with average human intelligence. These scenarios help give a sense of the flexibility of the forecast within IFs and how the AI index can be manipulated to better match user expectations.

⁴ A method of group decision-making and forecasting that involves successively gathering the opinions of experts to come to a consensus-style answer.

Table 8: Literature Survey on the Timeline for AGI Development		
Study	Details	Results – When will there be AGI?
Kurzweil (2005)	In his book the Singularity noted futurist Ray Kurzweil (now Google's Director of AI) laid out his forecast for the development of AGI.	AGI will be present around the year 2045.
Baum et al. (2011)	Assessment of expert opinion from participants at the AG-09 conference.	The consensus was that a large portion of the AI community believed AGI is possible around the middle of the current century.
Bostrom and Sandberg (2011)	Surveyed 35 participants at a human level intelligence conference in 2011.	Median results: 10% chance of AGI: 2028 50% chance of AGI: 2050 90% chance of AGI: 2150
Barrat and Goertzel (2011)	Surveyed participants at the AG-11 conference hosted by Google	Results: 42% of respondents: 2030 25% of respondents: 2050 20% of respondents: 2100 10% of respondents after 2100 2%: never
Muller and Bostrom (2014)	Electronic survey to hundreds of AI experts and researchers.	Median results: 10% chance of AGI: 2022 50% chance of AGI: 2040 90% chance of AGI: 2075

The scenarios displayed below underscore two fundamental uncertainties around the future of AI with respect to this forecasting exercise: (i) how “high” it can ultimately go (that is, what level can AI achieve); and (ii) how fast it will get there. The parameters added to IFs allow users to control both elements. The scenarios in Figure 8 both accelerate the pace of AI and affect its end level in 2100. Under Accelerated AI, the index reaches a score of close to 350 by 2100, whereas Stalled AI only achieves an index score of around 100 by 2100.


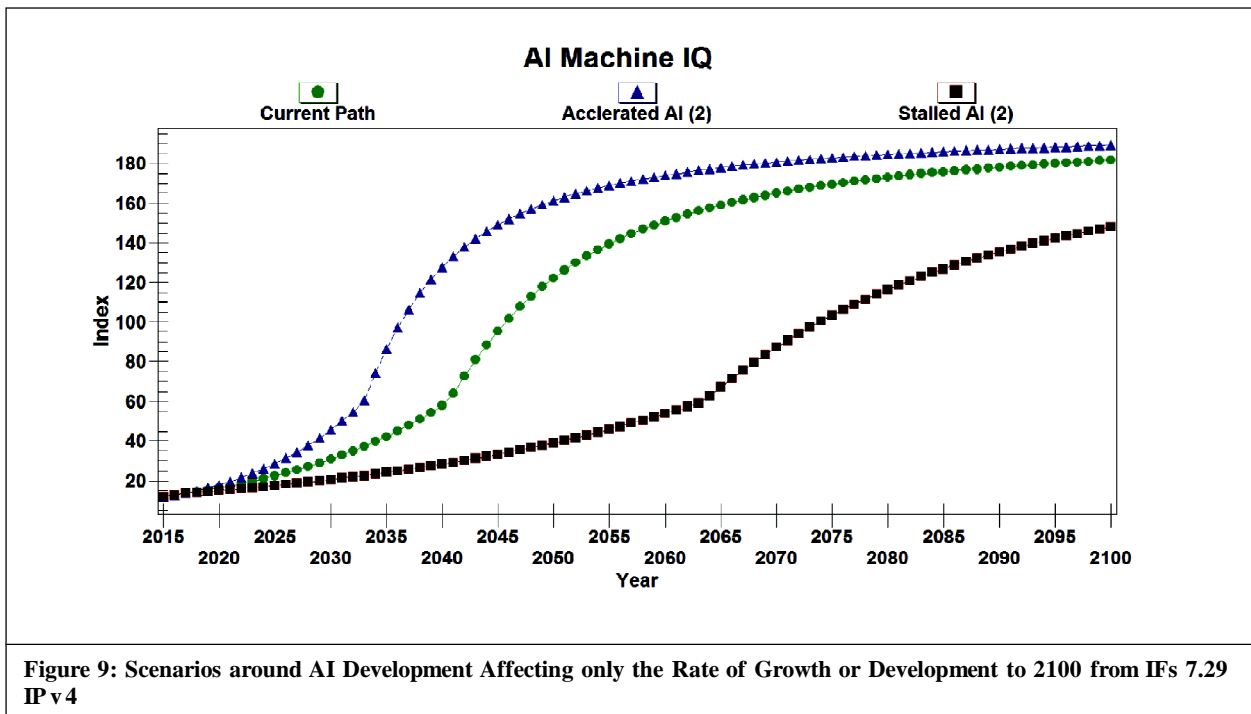


Figure 8: Scenarios Around AI Development Affecting both Rate of Growth and End Level in 2100 from IFs v. 7.29 IP 2

For the purpose of comparison and also to provide the readers with a sense of the customization built into the AI indices, Figure 9 displays the results of scenarios that affect the growth rate of AI technologies, but do not alter its end level by 2100. Both scenarios simulate a 50% increase or decrease in the rate of AI development relative to the current path. In Accelerated AI (2), AI converges towards an advanced machine IQ score of 180 more rapidly than in the current path. In this scenario we expect to see AGI emerge between 2035 and 2038, and superintelligent machines to come into being around mid-century. After 2050 AI technology growth slows as it converges towards a fixed level of superintelligence. In a similar pattern, Stalled AI (2) slows AI’s advance by 50% relative to the current path. In this scenario AI Machine IQ only begins to approach superintelligent levels by the end of century (approaching an index score of 150) but does not approach the maximum possible level of intelligence by the end of the horizon. AGI alone doesn’t emerge until mid-2060.



7. IFs: Exploring the Impacts of AI

As we have expressed throughout this report, AI will have deep impacts on many areas of human development. The utility of this quantitative forecast of AI development will be significantly enhanced by connecting the AI representation to other areas of the IFs model that would allow us to explore its impact at multiple levels over both the medium and long-term. The fact that IFs is integrated across so many different human development systems leaves it uniquely placed among other modeling efforts to capture the deep and wide-ranging impact of AI. Connecting AI to other areas of the model would have to be done through a set of carefully calibrated elasticity’s that could be freely adjusted by users. We propose to capture AI’s impact by on three areas in particular: economic productivity, labor, and international trade through production localization.

7.1. Economic Productivity

A near universal consensus in the literature suggests AI will improve economic productivity, but analysis on the depth of impact varies widely. Productivity, an assessment of output based on a fixed number of inputs, is a benchmark for the efficiency of production and technological progress (McGowan et al., 2015, p. 21). Nobel Prize winning economist Paul Krugman pointed out that with respect to economic growth, “productivity isn’t everything, but in the long run it is almost everything” (Krugman, 1994, p. 11). Fortunately, AI is poised to enhance productivity.

A 2016 report by Accenture, a consulting firm, laid out three avenues through which AI could enhance economic activity. The first is through intelligent automation, wherein AI is able to automate complex physical tasks, such as retrieving items in a warehouse. Increasingly intelligent AI machines are anticipated to be able to adapt across different tasks and industries. The second way AI will improve technology is by enhancing labor and capital, by freeing labor to act more creatively, imaginatively, and freely. The third way AI could enhance productivity is the

result of diffusion, whereby innovation catalyzed by AI moves through diverse sectors of the economy. For instance, driverless cars will not only fundamentally change how our automobiles work, they could entirely restructure the auto insurance industry, reduce traffic congestion, accidents, and associated hospital bills, and stimulate demand for smart infrastructure. The extent of the productivity increase in different sectors will be more closely tied to how susceptible each industry is to AI technologies and/or automation, rather than factors like the level of investment or the level of development of the country in question.

Most analysis of AI and productivity today focuses on estimating the benefits to productivity over the next decade or so. In 2015 Bank of America Merrill Lynch estimated that robots and AI technologies could bring add an estimated \$2 tn to US. GDP in efficiency gains over the next ten years, driven by the adoption of autonomous cars and drones. By their estimation robotics alone could drive productivity gains of 30% in many industries (Ma et al., 2015). A 2017 report from McKinsey Global Institute on labor and technology estimated that AI-driven automation could increase global productivity by 0.8% to 1.4% annually within the next few years. Accenture Consulting is even more optimistic, estimating that labor productivity will be between 11% and 37% higher in a sample of OECD countries in 2035 as a result of AI (Table 9).

Table 9: Forecasted Impacts of AI on Productivity in 2035	
Country	Percentage Increase in Labor Productivity in 2035 Compared to Base
Sweden	37%
Finland	36%
United States	35%
United Kingdom	25%
Belgium	17%
Spain	11%

Source: Purdy and Daugherty (2016)

Fewer attempts have been made to measure productivity and automation using historical data. One attempt by two researchers at Uppsala University and the London School of Economics used data from 1993 to 2007 in seventeen advanced economies. Across that period, the density of robots in manufacturing centers increased 150%, and both total factor productivity and wages increased. They find that robots increased GDP and labor productivity by 0.37 and 0.36 percentage points respectively. Although there is less research on automation and productivity using historical data, the argument for productivity gains from AI builds on a substantial body of evidence of productivity gains accruing to developed economies from the ICT boom in the 1990s and early 2000s. Research has identified positive productivity gains both within industries (Stiroh, 2002) and across countries and regions (Bloom et al., 2012; O'Mahony and Timmer, 2009; and Qiang, 2009).

Nevertheless, with respect to productivity, AI may be facing some strong headwinds. According to figures published in August 2016, US labor productivity levels declined for the third straight quarter (Azeez, 2016). This is symptomatic of broader trends in the US economy: between 2000 and 2007 annual productivity grew at around 2.6% between 2007 and 2016, it grew only by 1%. In the 1990's ICT gains helped US productivity grow by 2.2% per annum (Lam, 2017). This slowdown has not been restricted to just the US, nor is it necessarily specific to certain industries or sectors (Foda, 2016). Even by 2013, average productivity was 2% below levels seen prior to the 2008-2009 financial crisis across the OECD (McGowan et al., 2015). Declining productivity among advanced economies is a troubling phenomenon concerning to policymakers. A number of explanations have been put forth, including: (i) aging populations and structural economic inefficiencies (Gordon, 2012); (ii) labor reallocation challenges (Haltiwagner, 2011); (iii) increasingly bureaucratic and unwieldy firms (Hamel and Zanini, 2016); and (iv) slowing technology diffusion among firms and industries (McGowan et al., 2015).

A simpler explanation may be that technology has simply complicated calculations of GDP growth and productivity. Mainstream platforms from *The Economist* (2016b) to the World Economic Forum (Blanke, 2016) have recently catalogued issues with GDP as an indicator of economic growth. Mathematically, GDP represents the sum of all consumption, government spending and investment (plus exports minus imports). Governments commonly use GDP to set fixed growth targets. It provides a general picture of the health of a country’s economy.

The attachment to GDP has led to measures like GDP per capita representing proxies for standard of living economic wellbeing. And yet, economists increasingly point out that GDP is a poor indicator of economic and social wellbeing (Thompson, 2016). It says little about inclusive growth, or how the gains from growth are distributed. It says nothing about environmental degradation that may result from growth. It does not tell us whether growth is actually improving people’s lives. And yet, as David Pilling in an article for the *Financial Times* pointed out: “GDP may be anachronistic and misleading. It may fail entirely to capture the complex trade-offs between present




Figure 10: Internet Contribution to GDP

Source: Manyika and Roxburgh (2011)

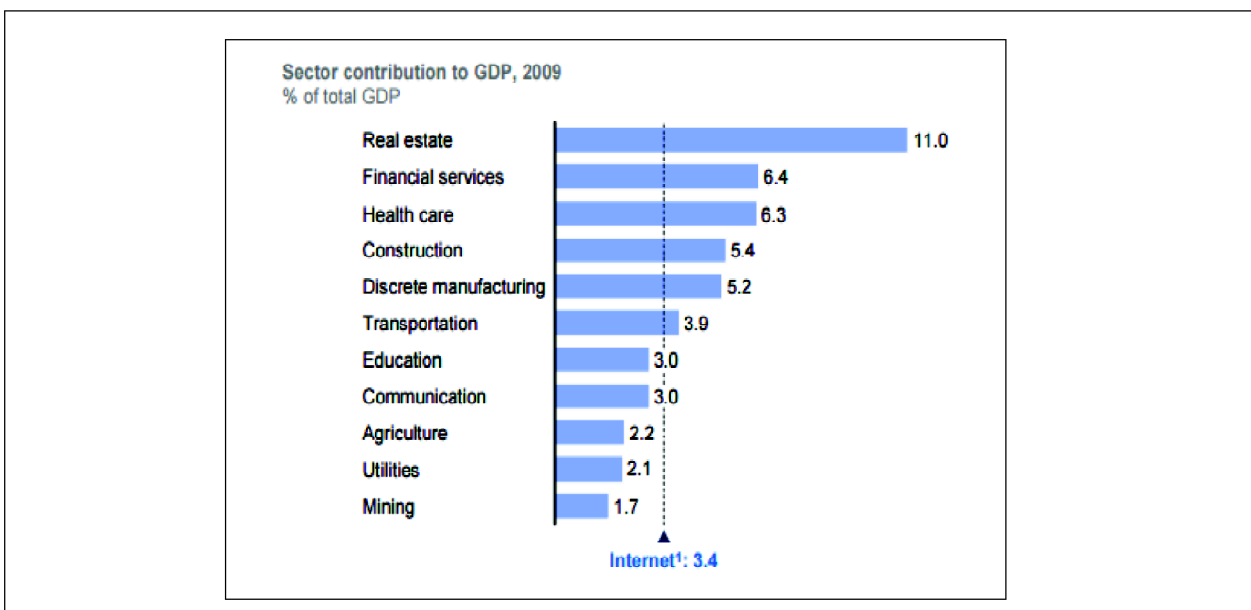


Figure 11: Sector Contribution to GDP

Source: Manyika and Roxburgh (2011)

and future, work and leisure, ‘good’ growth and ‘bad’ growth. Its great virtue, however, remains that it is a single, concrete number. For the time being, we may be stuck with it” (Pilling, 2014).

GDP is also problematic because it may not fully capture the benefits accruing from the digital economy. GDP has not kept pace with changes in the way the economy works (Libert and Beck, 2016). GDP misrepresents important activities related to things like knowledge creation, product quality improvements, stay-at-home parenting, or the gig economy. The sharing economy (think Uber or Airbnb) may not be properly valued through existing measurements. By one estimate, the sharing economy may have been worth around \$14 bn in 2014, and could grow to \$335 bn by 2025 (Yaraghi and Ravi, 2016). Misrepresenting or failing to capture such a rapidly growing industry would skew measurements of our true productivity.

With this debate over GDP and productivity in mind, any discussion over the impact of AI on productivity should entertain the concept of “consumer surplus,” that is the total value to the consumer for the use of an online good or service less any costs that consumers pay to access those services (Pélissié du Rausas et al., 2011). This has been advanced as an important concept in estimating the value of the digital economy.

A 2011 report from McKinsey Global Institute put the value of the “internet economy” at around \$8 tn, accounting for more than 3% of global GDP among developed countries (Figure 10).⁵ If it were a sector, the internet would be more significant than agriculture or utilities (Figure 11). Across the different countries explored in the report, the total consumer surplus ranged from \$10 bn in Germany and France to near \$64 bn in the United States. A separate but related piece of McKinsey analysis looked at the economic value of internet searching in five major economies (Brazil, France, India, Germany, and the United States). They estimated internet search was worth close to \$870 bn across the global economy. Of that, roughly 31% (\$240 bn) is not captured in GDP statistics, but represents a consumer surplus, or the value accruing from the benefits of convenience, lower prices, and ease of information access (Bughin et al., 2011).

Other studies have attempted to measure the impact of the internet on GDP and consumer surplus. One 2009 study completed by consultants with Harvard Business School estimated that approximately 2% of Americans were employed directly or indirectly by internet-related activities (advertising, commerce, IT infrastructure, maintenance), generating close to \$300 bn in wages. In addition to jobs, the internet adds an estimated \$175 bn to the US economy through retail, advertising, and payments to internet service providers. Moreover, between work and leisure, they estimated Americans spend close to 68 hours per month on the internet, which produces an estimated \$680 bn in value (Quelch, 2009). A 2016 study from Georgetown University estimated that for every \$1 spent using Uber, a US-based ride-sharing service, \$1.60 of consumer surplus was generated. They estimated that across the US, Uber helped generate \$6.8 bn in consumer benefits (Cohen et al., 2016).

Nevertheless, consumer surplus is notoriously difficult to measure. Measuring surplus requires knowing the demand for a product. But many digital services like Facebook and Google are free. Without a price, it is difficult to quantify demand. Moreover, users of digital services like Facebook derive different levels of surplus or satisfaction. The value someone places on Facebook is dependent on their networks; if more of their friends are active on Facebook and social media, they will derive greater value. These kinds of implications raise questions about whether it is possible to derive a single demand curve for digital products. At the same time, the growth of the internet and the digital economy is undeniable, and many of its welfare-producing activities are not currently well captured in GDP measurements. New methods of capturing value-add in the digital age will produce a more accurate picture of productivity, particularly in the developed world, and allow researchers and policymakers to respond and adapt appropriately.

7.2. Labor

At present, there is little that captures the attention of mainstream media and policymakers like the potential impact of AI on labor, particularly through the computerization and automation of jobs. At the 2017 World Economic Forum in Davos, a panel of technology leaders and AI experts focused not on the potential for large profits and the business gains, but how to deal with those left behind in the digital age (Bradshaw, 2017). The populist backlash to the impacts of globalization that culminated in Brexit and the election of Donald Trump as US President, coupled with the rise of populist parties in Europe shows that these concerns are well founded and can have real political implications. Adding fuel to the flames of populist sentiments are headline-grabbing analyses such as the

⁵ Based on an analysis of 13 economies accounting for 70% of global GDP.

2013 report by from Oxford University that estimated close to 47% of jobs in the US labor market were at risk of automation in the next 10 years (Frey and Osborne, 2013). Perhaps AI is leading us all into a jobless future.

In reality, it is difficult to quantify the effect of technology on labor, and even more difficult to predict the scope and breadth of future automation. For every headline predicting massive social dislocation from AI, there are often corresponding analyses predicting that AI will unleash a new wave of jobs in new industries emerging from the AI revolution. The optimists argue that AI will take over jobs that are dull and dangerous, freeing up human labor for more creative and fulfilling tasks. This remains a widely debated and hotly contested issue. Let us look at some of the forecasted implications.

The 2016 World Economic Forum produced a background report on the future of jobs. In the report, they surveyed 15 of the world's largest economies, comprising approximately 1.86 billion workers or 65% of the total global workforce. They concluded that AI will lead to a net loss of 5.1 million jobs between 2015 and 2020 (7.2 million lost, 2.1 million gained) (Schwab and Samans, 2016b). McKinsey Global Institute estimated that activities accounting for close to \$15 trillion in wages globally could be automated by adapting current technologies, and that half of all work today could be automated away by 2055 (Manyika et al., 2017). While developed countries are likely to experience the effects of AI more rapidly because their economies depend more on technology, the effects are by no means restricted to the developed world. According to the World Bank, as many as 77% of jobs in China, 69% in India, and 85% in Ethiopia may be at risk of automation (World Bank Group, 2016). The jobs at risk for automation are highly repetitive tasks in structured environments, and data collection and analysis. Laborers in developing countries may also be sensing a trend: according to a survey of workers in 13 countries, 80% of respondents in China and 62% in India felt AI would replace human labor in repetitive tasks. In Germany and the UK by contrast, only 39 and 45% of respondents felt the same way (Wong, 2016). The jobs at risk for automation are highly repetitive tasks in structured environments, and data collection and analysis. The sectors most at-risk in the US market include manufacturing, food service, retail, and some service sectors (Manyika et al., 2017).

Estimating the impact of AI on labor also forces us to think about jobs as a series of tasks rather than as one monolithic entity. The same McKinsey Global Institute Report actually estimated that only 5% of jobs could be fully automated, but that close to 60% of jobs in the US market could be up to 30% automated at the task level within the next 20 years. This adds weight to the argument of optimists that AI will actually free up human labor for more meaningful activities. A 2016 report from the OECD looked at the prospects of automation across OECD countries. Employing similar estimation techniques as the Oxford paper but controlling for within-job tasking, they estimated the risk of computerization and found, on average, nine percent of jobs are at-risk (Arntz et al., 2016).

There is more evidence that technology creates jobs by creating new products, changing preferences, and inducing competitiveness. In a 2016 report, analysts from Deloitte looked at the history of jobs and technology in the US and UK between 1871 and today. They concluded that over the past 144 years, technology has created more jobs than it has cost. While technology has replaced some jobs, it has created new ones in knowledge and service sectors like medicine and law. Technology has reduced the cost of basic goods and raised incomes, prompting the creation of new jobs to meet changing demand patterns (Stewart et al., 2015).


7.3. Localization of Production and International Trade

Another trend that could be significantly impacted by the rise of AI deserves consideration: reshoring and the localization of production. Automated technologies are making it increasingly inexpensive for companies to produce goods at home, reducing the need for offshoring in search of cheap labor and competitive prices. In the US there has been discussion around the idea of reshoring and anecdotal evidence suggests it is happening, yet critics contest the US government does not maintain exhaustive data on reshoring and that the definition of reshoring itself remains contested, thus it is difficult to say whether it represents an industry-wide trend (Rivkin, 2014).

There is anecdotal evidence to hint at a trend. The term reshoring refers to the process of relocating production centers in typically developed countries. A 2012 MIT survey of 340 participants from the manufacturing industry found that 33% were "considering" bringing manufacturing back to US shores (Simchi-Levit, 2012), while a 2013 report in the Economist found that between 37 and 48% of manufacturing firms with \$1 bn or more in revenues were considering reshoring or had already begun the process (The Economist, 2013). Individual examples of large companies moving production back to the US or Europe have appeared in the media frequently in recent years (Oldenski, 2015). For instance:

- In 2009 General Electric relocated production of water heaters from China to Kentucky;
- In 2010 Master Lock returned 100 jobs to Milwaukee, Wisconsin;
- In 2012 Caterpillar opened a new plant in Texas;
- In 2014 General Motors moved a production plant from Mexico to Tennessee;
- In 2015 Ford began announced it would begin producing engines at its Cleveland auto plant; and
- In August 2016, Adidas opened its first manufacturing plant in Germany in over 30 years.

The anecdotal evidence does not necessarily constitute a trend. For instance, the “reshoring index,” put together by consultancy group ATKearney reports that there were only about 60 cases of reshoring in the US in 2015, down from 208 cases in 2014 (Figure 12). The index estimates that there were 210 cases in 2013, 104 in 2012, and 64 in 2011, small figures when considering that US multinational corporations employ as many as 36 million people worldwide (Oldenski, 2015). These examples of reshoring also say nothing of any concurrent offshoring activity that may have happened during the same period.



Nevertheless, the fact remains that automation, coupled with low-cost energy and rising wages in the developing world, particularly China and India, has the potential to make companies rethink where they base their operations. There are also incentives for companies to base their operations close to their primary markets to reduce shipping time and costs, and improve their ability to respond to local market needs and fluctuations. Moreover, in today’s populist political climate, there are incentives that encourage companies to invest locally. In an AI-led world, it is possible that a majority of production happens locally, reducing the necessity for the cross-border movement of goods and services.

The energy sector is one where this potential trend could manifest itself with significant implications for global trade. AI has the potential to disrupt current energy patterns by driving growth in renewable production that causes a reduction in the volume of international trade in traditional energy products, particularly fossil fuels.

AI is already improving the efficacy of renewable energy production. A core challenge in harnessing renewable energies like wind and solar is their intermittency. Machine learning is helping to overcome this hurdle by crunching real-time data on weather conditions to produce accurate forecasts, allowing companies to better harness these sources (Bullis, 2014). In Germany, companies are using machine learning to crunch data and predict wind generation capacity in 48 hour increments which allows the national energy grid to respond to energy demand without relying on traditional energy sources to cover shortfalls (Thompson, 2016).

AI is also poised to boost renewable generation by significantly enhancing demand-side efficiency. Machine learning, coupled with smart meters and smart applications, can help large grid systems identify consumption patterns and adjust energy provision and storage accordingly. AI technology is being applied to mine data that allows grid systems to come up with suitable and appropriate risk/reward mechanisms that both incentivize their customers to participate in smart energy and obtain measurable benefits (Robu, 2017). We can already see some of these patterns beginning to emerge. For instance, 2016 was the cleanest year on record for the UK, where coal-fired energy production fell to under 10% of total production, down from 40% in 2012. Wind power generation alone was higher than coal, at 10.2% (Wilson and Staffell, 2017). On a Sunday in May 2016, close to 100% of Germany's power demand was met using only renewable sources, primarily wind and solar. For a short 15 minute window during that day, power prices in Germany actually went negative (Shankelman, 2016).

The growth of renewable energy capable of being domestically sourced and harnessed has important implications for global trade. Crude oil and its derivatives remain the most valuable traded commodity in the world. According to the UN Conference on Trade and Development (UNCTAD), trade in oil, gas, and petroleum products were estimated at between \$1 and \$2 tn in 2014 and 2015, among the largest of the 25 categories of goods and services tracked by the organization. British Petroleum (BP) estimated that in 2015 close to 1.02 billion tons of crude oil were exported and 1.9 billion tons were imported (British Petroleum, 2016). The global energy trade remains significant today, but renewable generation could slow that trade. The IFs Current Path Forecast estimates that by 2050 close to 40% of the world's energy production will come from renewable sources, up from around 6% today.

8. Conclusion

This report has outlined a conceptual framework, operationalization, and forecast of AI to 2100. It has also laid out the potential to model the impact of AI within IFs with a particular focus on economic productivity, labor, and international trade and production localization. We will not try to summarize our findings here but instead encourage the reader to revisit the executive summary. We conclude this report by reminding readers of the benefits that quantitative modeling can bring to the understanding of AI and its disparate impacts. We have been forthcoming about the level of uncertainty surrounding this forecasting exercise and have designed the AI representation to provide maximum user flexibility and freedom. AI development is rapidly unfolding and is expected to have broad social and global impact. To better unpack and understand the future impact of AI requires connecting the AI forecast representation to other areas of the IFs model. This is an important area of future research and IFs remains uniquely placed to pursue this endeavor. We fully believe further exploration and forecasting of these issues will be beneficial to the research community and broader public alike.

References

- Ackerman, E. (2016). *Winograd Schema Challenge Results: AI Common Sense Still a Problem, for Now*. IEE Spectrum. Available: <http://spectrum.ieee.org/automaton/robotics/artificial-intelligence/winograd-schema-challenge-results-ai-common-sense-still-a-problem-for-now>. (accessed on 26 April 2017).
- Adams, S.S., Arel, I., Bach, J., Coop, R., Furlan, R., Goertzel, B., Hall, J.S., Samsonovich, A., Scheutz, M., Schlesinger, M., Shapiro, S.C., and Sowa, John. (2012). *Mapping the Landscape of Human-Level Artificial Intelligence*. *AI Magazine* 33, 25-42.
- AI Impacts. (2016). *Global Computing Capacity*. *AI Impacts*. 16 February 2016. Available: <http://aiimpacts.org/global-computing-capacity/> (accessed on April 9, 2017).
- Albergotti, R. (2014). *Zuckerberg, Musk Invest in Artificial-Intelligence Company*. *Wall Street Journal*. 21 March 2014. Available: <https://blogs.wsj.com/digits/2014/03/21/zuckerberg-musk-invest-in-artificial-intelligence-company-vicarious/> (accessed on August 21, 2016).
- Amazon. (2017). *Amazon Machine Learning*. Available: <https://aws.amazon.com/machine-learning/> (accessed on April 13, 2017).
- Arntz, M., Gregory, T., and Zierahn, U. (2016). *The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis*. OECD Social, Employment, and Migration Working Papers, 189. Paris: OECD.

- AT Kearney. (2015). *U.S. Reshoring: Over Before It Began?* AT Kearney: U.S. Reshoring Index. Available: <https://www.atkearney.com/documents/10192/7070019/US+Reshoring.pdf/14777afa-0c14-460e-937b-11233be340b6>. (accessed on April 7, 2017).
- Azeez, W. (2016). *The Morning Ledger: Productivity Slump Threatens Long-Term U.S. Growth.* *Wall Street Journal*. 10 August 2016. Available: <https://blogs.wsj.com/cfo/2016/08/10/the-morning-ledger-productivity-slump-threatens-long-term-u-s-growth/> (accessed on April 12, 2017).
- Barrat, J., and Goertzel, B. (2011). *How Long Till AGI? - Views of AGI-11 Conference Participants.* *Humanity+ Magazine*. 16 September 2011. Available: <http://hplusmagazine.com/2011/09/16/how-long-till-agi-views-of-agi-11-conference-participants/> (accessed on 2 April 20, 2017).
- Baum, S.D., Goertzel, B., and Goertzel, T.G. (2011). *How Long Until Human-Level AI? Results from an Expert Assessment.* *Technological Forecasting and Social Change* 78, 185-195.
- Blanke, J. (2016). *What is GDP, and how are we misusing it?* World Economic Forum. April 13, 2016. Available: <https://www.weforum.org/agenda/2016/04/what-is-gdp-and-how-are-we-misusing-it/> (accessed on July 31, 2016).
- Bloom, N., Sadun, R., and Reenen, J.V. (2012). *Americans Do IT Better: US Multinationals and the Productivity Miracle.* *American Economic Review*, 102, 167-201.
- Bostrom, N. (1998). *How Long Before Superintelligence?* *International Journal of Future Studies*, 2.
- Bostrom, N. (2014). *Superintelligence: Paths, Dangers, Strategies*, First edition. (Oxford: Oxford University Press).
- Bostrom, N., and Sandberg, A. (2011). *Machine Intelligence Survey*, Technical Report #2011-1, Future of Humanity Institute, Oxford University: pp. 1-12.
- Bourzac, K. (2016). *Intel: Chips Will Have to Sacrifice Speed Gains for Energy Savings.* *MIT Technology Review*. 05 February 2016. Available: <https://www.technologyreview.com/s/600716/intel-chips-will-have-to-sacrifice-speed-gains-for-energy-savings/> (accessed on February 21, 2017).
- Bradshaw, T. (2017). *Tech Leaders at Davos fret over effect of AI on jobs.* *Financial Times*. January 29, 2017. Available: <https://www.ft.com/content/744ad7fa-de66-11e6-9d7c-be108f1c1dce> (accessed on March 3, 2017).
- British Petroleum. (2016). *Statistical Review of World Energy*. June 2016. BP. Available: <http://www.bp.com/en/global/corporate/energy-economics/statistical-review-of-world-energy/oil/oil-trade-movements.html> (accessed on March 18, 2017).
- Bughin, J., Corb, L., Manyika, J., Nottebohm, O., Chui, M., Barbat, B. de M., and Said, R. (2011). *The Impact of Internet Technologies: Search.* McKinsey Global Institute. August 2011. Available: <http://www.mckinsey.com/business-functions/marketing-and-sales/our-insights/measuring-the-value-of-search> (accessed on April 17, 2017).
- Bullis, K. (2014). *Smart Wind and Solar Power.* *MIT Technology Review*. May/June 2014. Available: <https://www.technologyreview.com/s/526541/smart-wind-and-solar-power/> (accessed on September 9, 2016).
- Byrnes, N. (2016). *AI Hits the Mainstream.* *MIT Technology Review*. March 28, 2016. Available: <https://www.technologyreview.com/s/600986/ai-hits-the-mainstream/> (accessed on April 9, 2017).
- CB Insights. (2017). *The 2016 AI Recap: Startups See Record High in Deals and Funding.* CB Insights. Available: <https://www.cbinsights.com/blog/artificial-intelligence-startup-funding/> (accessed on April 12, 2017).
- Chen, N., Christensen, L., Gallagher, K., Mate, R., and Rafert, G. (2016). *Global Economic Impacts Associated with Artificial Intelligence.* *The Analysis Group*. Available: http://www.analysisgroup.com/uploadedfiles/content/insights/publishing/ag_full_report_economic_impact_of_ai.pdf (accessed on August 12, 2016).
- Cisco. (2016). *Cisco Global Cloud Index White Paper: 2015-2020.* Cisco. Available: <http://www.cisco.com/c/dam/en/us/solutions/collateral/service-provider/global-cloud-index-gci/white-paper-c11-738085.pdf> (accessed on April 13, 2017).

- Cohen, P., Hahn, R., Hall, J., Levitt, S., Metcalfe, R. (2016). *Using Data to Estimate Consumer Surplus: The Case of Uber*. Working Paper No. 22627. National Bureau of Economic Research, Cambridge, Mass.
- DARPA. (2017). *Urban Challenge*. archived. Available: <http://archive.darpa.mil/grandchallenge/> (accessed on 15 March 2017).
- Diamandi, Z., Dubey, A., Pleasance, D., and Vora, A. (2011). *Winning in the SMB Cloud: Charting a Path to Success*. McKinsey & Company, New York, NY.
- Economist Intelligence Unit. (2016). *Ascending Cloud: The Adoption of Cloud Computing in Five Industries*. *The Economist*. March 1, 2016. Available: <https://www.eiuperspectives.economist.com/technology-innovation/ascending-cloud-adoption-cloud-computing-five-industries-0> (accessed on April 4, 2017).
- Ericsson. (2016). *Mobility Report: 2016. On the Pulse of the Networked Society*. Ericsson. Available: <https://www.ericsson.com/res/docs/2016/ericsson-mobility-report-2016.pdf> (accessed on April 8, 2017).
- Evans, T.G. (1964). *A Heuristic Program to Solve Geometric-Analogy Problems*, in: *Spring Joint Computer Conference*, 328-338.
- Ferucci, D., Brown E., Chu-Carroll, J., Fan, J., Gondek, D., Kalyanpur AA., et al. (2010). *Building Watson: An Overview of the DeepQA Project*. *AI Magazine*. Fall. Available: <http://www.aaai.org/Magazine/Watson/watson.php> (accessed on March 3, 2017).
- Fischetti, M. (2011). *Computers versus Brains*. *Scientific American*. November 1, 2011. Available: <https://www.scientificamerican.com/article/computers-vs-brains/> (accessed on June 8, 2016).
- Foda, K. (2016). *The Productivity Slump: A Summary of the Evidence*. Brookings Institute. Available: <https://www.brookings.edu/research/the-productivity-slump-a-summary-of-the-evidence/> (accessed on April 18, 2017).
- French, R.M. (2000). *The Turing Test: The First 50 years*. *Trends in Cognitive Sciences*, 4, 115-122.
- Frey, C.B., and Osborne, M.A. (2013). *The Future of Employment: How Susceptible are Jobs to Computerization?* September 17, 2013. Oxford Martin School, Oxford: UK.
- Frey, C.B., Osborne, M.A., Holmes, C., Rahbari, E., Curmi, E., Garlick, R., Chua, J., Friedlander, G., Chalif, P., and McDonald, G. (2016). *Technology at Work V2.0. Citi GPS*. January 3, 2016. Oxford Martin School & Citi. Available: http://www.oxfordmartin.ox.ac.uk/downloads/reports/Citi_GPS_Technology_Work_2.pdf (accessed on July 9, 2016).
- Ginsberg, M.L. (1993). *Essentials of Artificial Intelligence*. (San Mateo: Morgan Kaufmann Publishers).
- Goertzel, B. (2012). *What Counts as a Conscious Thinking Machine?* *NewScientist*. September 5, 2012. <https://www.newscientist.com/article/mg21528813-600-what-counts-as-a-conscious-thinking-machine/> (accessed on April 19, 2017).
- Gordon, Robert, J. (2012). *Is US Economic Growth Over? Faltering Innovation Confronts the Six Headwinds*. Working Paper No. 18315. National Bureau of Economic Research. Cambridge, Mass.
- Guszcza, J., Lucker, J., and Lewis, H. (2013). *Too Big to Ignore*. Deloitte University Press. January 31, 2013. Available: <https://dupress.deloitte.com/dup-us-en/deloitte-review/issue-12/too-big-to-ignore.html> (accessed on February 22, 2017).
- Haltiwagner, J. (2011). *Firm Dynamics and Productivity Growth*. *European Investment Bank Papers*, 16, 116-136.
- Hamel, G., and Zanini, M. (2016). *More of Us Are Working in Big Bureaucratic Organizations than Ever Before*. *Harvard Business Review*. July 5, 2016. Available: https://hbr.org/2016/07/more-of-us-are-working-in-big-bureaucratic-organizations-than-ever-before?utm_source=twitter&utm_medium=social&utm_campaign=harvardbiz (accessed on August 18, 2016).
- Harnad, S. (1990). *The Symbol Grounding Problem*. *Physica D: Nonlinear Phenomena*, 42, 335-346.
- Hawkins, J., and Dubinsky, D. (2016). *What Is Machine Intelligence Vs. Machine Learning Vs. Deep Learning Vs. Artificial Intelligence (AI)?* *Numenta*. January 11, 2016. Available: <http://numenta.com.s3-website-us-west->

- 2.amazonaws.com/blog/2016/01/11/machine-intelligence-machine-learning-deep-learning-artificial-intelligence/ (accessed on July 7, 2016).
- Hernández-Orallo, J. (2017). *The Measure of All Minds: Evaluating Natural and Artificial Intelligence*. Cambridge: Cambridge University Press, United Kingdom.
- Hilbert, M., and Lopez, P. (2012). How to Measure the World's Technological Capacity to Communicate, Store and Compute Information Part I: Results and Scope. *International Journal of Communication*. 6, 956-979.
- Hof, R. (2013). Deep Learning. *MIT Technology Review*. Available: <https://www.technologyreview.com/s/513696/deep-learning/> (accessed on July 10, 2016).
- Hollinger, P. (2016). Meet the Cobots: Humans and Robots Together on the Factory Floor. *Financial Times*. May 4, 2016. Available: <https://www.ft.com/content/6d5d609e-02e2-11e6-af1d-c47326021344> (accessed on 13 July 2016).
- Howard, P.N. (2015). Sketching out the Internet of Things Trendline. Brookings Institute. June 9, 2015. Available: <https://www.brookings.edu/blog/techtank/2015/06/09/sketching-out-the-internet-of-things-trendline/> (accessed on March 3, 2017).
- Hughes, B.B. (2019). *International Futures: Building and Using Global Models*. London: Elsevier Academic Press. ISBN: 9780128042717.
- Hughes, Neil. (2014). Tests find Apple's Siri improving, but Google Now Voice Search Slightly Better. *Appleinsider*. July 22, 2014. Available: <http://appleinsider.com/articles/14/07/22/tests-find-apples-siri-improving-but-google-now-voice-search-slightly-better> (accessed on March 24, 2017).
- IDC. (2016). Worldwide Public Cloud Services Spending Forecast to Reach \$195 Billion by 2020, According to IDC. International Data Corporation. August 10, 2016. Available: <http://www.idc.com/getdoc.jsp?containerId=prUS41669516> (accessed on March 12, 2017).
- International Federation of Robotics. (2016). World Robotics 2016 Industrial Robots. *International Federation of Robotics*. Available: https://ifr.org/img/uploads/Executive_Summary_WR_Industrial_Robots_20161.pdf (accessed on April 20, 2017).
- International Federation of Robotics. (2017). History of Robotics. *International Federation of Robotics*. Available: <https://ifr.org/robot-history>. (accessed on March 3, 2017).
- International Telecommunications Union. (2016). ICT Facts and Figures 2016. *International Telecommunications Union*. Available: <http://www.itu.int/en/ITU-D/Statistics/Pages/facts/default.aspx> (accessed on April 19, 2017).
- Isaac, M., and Boudette, N.E. (2017). Ford to Invest \$1 Billion in Artificial Intelligence Start-Up. *The New York Times*. February 10, 2017. Available: <https://www.nytimes.com/2017/02/10/technology/ford-invests-billion-artificial-intelligence.html> (accessed on March 5, 2017).
- Knight, W. (2015). Baidu System Rivals People at Speech Recognition. *MIT Technology Review*. December 16, 2015. Available: <https://www.technologyreview.com/s/544651/baidus-deep-learning-system-rivals-people-at-speech-recognition/> (accessed on March 17, 2017).
- Knight, W. (2017). Reinforcement Learning. *MIT Technology Review*. Available: <https://www.technologyreview.com/s/603501/10-breakthrough-technologies-2017-reinforcement-learning/> (accessed on March 7, 2017).
- Krugman, P.R. (1994). *The Age of Diminished Expectations: US Economic Policy in the 1990s*, 3rd Edition. Cambridge: MIT Publishing.
- Kurzweil, R. (2005). *The Singularity is Near: When Humans Transcend Biology*. Viking: New York.
- Lam, B. (2017). Can Declining Productivity Growth Be Reversed? *The Atlantic*. March 14, 2017. Available: <https://www.theatlantic.com/business/archive/2017/03/productivity-interest-rate/519522/> (accessed on April 20, 2017).

- Le, Q., Ranzato, M.A., Monga R., Devin, M., Chen K., Corrado, G.S., et al. (2012). *Building High-Level Features Using Large Scale Unsupervised Learning*. In *International Conference on Machine Learning*. Edinburgh, UK.
- Lewis-kraus, G. (2016). *The Great A.I. Awakening*. *The New York Times*. December 14, 2016. Available: <https://www.nytimes.com/2016/12/14/magazine/the-great-ai-awakening.html> (accessed on March 7, 2017).
- Libert, B., and Beck, M. (2016). *GDP is a Widely Flawed Measure for the Digital Age*. *Harvard Business Review*. July 28, 2016. Available: <https://hbr.org/2016/07/gdp-is-a-wildly-flawed-measure-for-the-digital-age> (accessed on March 20, 2017).
- Ma, B., Nahal, S., and Tran, F. (2015). *Robot Revolution- Global Robot & AI Primer*. Bank of America Merrill Lynch. December, 1-13.
- Manyika, J., and Roxburgh, C. (2011). *The Great Transformer: The Impact of the Internet on Economic Growth and Prosperity*. McKinsey Global Institute. New York.
- Manyika, J., Chui, M., Bughin, J., Dobbs, R., Bisson, P., and Marrs, A. (2013). *Disruptive Technologies: Advances that will Transform Life, Business, and the Global Economy*. McKinsey Global Institute. New York.
- Manyika, J., Chui, M., Bisson, P., Woetzel, J., Dobbs, R., Bughin, J., and Aharon, D. (2015). *Unlocking the Potential of the Internet of Things*. McKinsey Global Institute, New York.
- Manyika, J., Chui, M., Miremadi, M., Bughin, J., George, K., Willmott, P., and Dewhurst, M. (2017). *Harnessing Automation for a Future that Works*. McKinsey Global Institute. New York.
- Markets and Markets. (2016a). *Machine Learning as a Service (MLaaS)- Global Forecast to 2021*. *Markets and Markets*. November 2016. Available: <http://www.marketsandmarkets.com/Market-Reports/machine-learning-as-a-service-market-183667795.html> (accessed on April 23, 2016).
- Markets and Markets. (2016b). *Natural Language Processing Market worth 16.07 billion by 2021*. *Markets and Markets*. Available: <http://www.marketsandmarkets.com/PressReleases/natural-language-processing-nlp.asp> (accessed on April 20, 2017).
- McCorduck, P. (2004). *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence, 25th Anniversary Update*. (Maine: A.K. Peters).
- McGowan, M.A., Andrews, D., Criscuolo, C., and Nicoletti, G. (2015). *The Future of Productivity*. OECD, Paris, France.
- Metz, R. (2013). *AI Startup Says it has Defeated Captchas*. *MIT Technology Review*. October 28, 2013. Available: <https://www.technologyreview.com/s/520581/ai-startup-says-it-has-defeated-captchas/> (accessed on February 28, 2017).
- Michigan News*. (2015). *U-M Opens Mcity Test Environment for Connected and Driverless Vehicles*. July 20, 2015. Available: <http://ns.umich.edu/new/multimedia/videos/23020-u-m-opens-mcity-test-environment-for-connected-and-driverless-vehicles> (accessed on March 20, 2017).
- Microsoft. (2016). *Microsoft Expands Artificial Intelligence (AI) Efforts with Creation of New Microsoft AI and Research Group [Press Release]*. 29 September. <https://news.microsoft.com/2016/09/29/microsoft-expands-artificial-intelligence-ai-efforts-with-creation-of-new-microsoft-ai-and-research-group/> (accessed 18 January 2022).
- Moon, P. (2007). *Three Minutes with Steve Wozniak*. *PCWorld*. July 19, 2007. Available: <http://www.pcworld.com/article/134826/article.html> (April 28, 2017).
- Moore, G.E. (1965). *Cramming More Components onto Integrated Circuits*. *Electronics*, 8, 114-117.
- Moosavi-Dezfooli, S.-M., Fawzi, A., Fawzi, O., and Frossard, P. (2016). *Universal Adversarial Perturbations*, in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- Muehlhauser, L. (2013). *What is AGI?* Machine Intelligence Research Institute. August 11, 2013. Available: <https://intelligence.org/2013/08/11/what-is-agi/> (accessed on April 1, 2017).

- Muller, V.C., and Bostrom, N. (2014). Future Progress in Artificial Intelligence: A Survey of Expert Opinon, in: Vincent C. Muller (ed.), *Fundamental Issues of Artificial Intelligence*. Synthese Library; Berlin: Springer.
- Newell, A., Shaw, J.C., and Simon, H.A. (1959). Report on A General Problem-Solving Program No. P-1584. Rand Corporation. Washington, DC.
- Ng, A. (2016). Andrew Ng: What AI Can and Can't Do. *Harvard Business Review*. November 9, 2016. Available: <https://hbr.org/2016/11/what-artificial-intelligence-can-and-cant-do-right-now> (accessed on March 17, 2016).
- Nguyen, A., Yosinki, J., and Clune, J. (2015). Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images. *In Computer Vision and Pattern Recognition, IEEE*, 2015.
- National Institute of Standards and Technology. (2012). NIST Open Machine Translation Evaluation (OpenMT12). NIST. August 28, 2012. Available: <https://www.nist.gov/multimodal-information-group/openmt12-evaluation-results#progress> (accessed on March 16, 2017).
- Nordhaus, W.D. (2001). *The Progress of Computing*. Version 4.4. Yale University and National Bureau of Economic Research. August 2001.
- Oldenski, L. (2015). Reshoring by US Firms: What Do the Data Say? Policy Brief No. PB15-14. Peterson Institute for International Economics. Available: <https://piie.com/publications/policy-briefs/reshoring-us-firms-what-do-data-say> (April 21, 2017).
- O'Mahony, M., and Timmer, M.P. (2009). Output, Input and Productivity Measures at the Industry Level: The EU KLEMS Database. *The Economic Journal*, 119, F374-F403.
- Parloff, R. (2016). Why Deep Learning Is Suddenly Changing Your Life. *Fortune*. September 28, 2016. Available: <http://fortune.com/ai-artificial-intelligence-deep-machine-learning/> (accessed on April 12, 2017).
- Pattani, A. (2016). Building the City of the Future- at a \$41 trillion Price Tag. CNBC. Available: <http://www.cnbc.com/2016/10/25/spending-on-smart-cities-around-the-world-could-reach-41-trillion.html> (accessed on April 21, 2017).
- Pélessié du Rausas, M., Manyika, J., Hazan, E., Bughin, J., Chui, M., and Said, R. (2011). *Internet Matters: The Net's Sweeping Impact on Growth, Jobs, and Prosperity*. McKinsey Global Institute, New York.
- Pilling, D. (2014). Has GDP Outgrown its Use? *Financial Times*. July 4, 2014. Available: <https://www.ft.com/content/dd2ec158-023d-11e4-ab5b-00144feab7de> (March 12, 2016).
- Purdy, M., and Daugherty, P. (2016). Why Artificial Intelligence is the Future of Growth. *Accenture*. Available: <https://www.accenture.com/us-en/insight-artificial-intelligence-future-growth> (accessed on April 20, 2017).
- Qiang, C.Z.-W. (2009). *Telecommunications and Economic Growth* (Unpublished Paper). World Bank, Washington, DC.
- Quelch, J. (2009). Quantifying the Economic Impact of the Internet. *Harvard Business School*. Available: <http://hbswk.hbs.edu/item/quantifying-the-economic-impact-of-the-internet> (accessed on April 22, 2017).
- Regalado, A. (2014). Is Google Cornering the Market on Deep Learning? *MIT Technology Review*. 29 June 29, 2014. Available: <https://www.technologyreview.com/s/524026/is-google-cornering-the-market-on-deep-learning/> (accessed on May 6, 2016).
- Rivkin, J. (2014). Reshoring US Manufacturing. *Economist Intelligence Unit*. December 15, 2014. Available: <https://www.eiuperspectives.economist.com/economic-development/reshoring-us-manufacturing> (accessed on April 20, 2017).
- Robu, V. (2017). How AI Can Help Renewables Work Better for the Energy Grid. *Newsweek*. January 30, 2017. Available: <http://www.newsweek.com/uk-grid-artificial-intelligence-energy-supply-supply-and-demand-550148> (accessed on April 29, 2017).
- Russakovsky, O., Deng J., Hao S., Krause J., Satheesh, S., Ma, S., et al. (2015). ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115, 211-252.
- Rutkin, A.H. (2017). The Tiny Changes that can Cause AI to Fail. *BBC FutureNow*. April 2017. Available: <http://www.bbc.com/future/story/20170410-how-to-fool-artificial-intelligence> (accessed on May 1, 2017).

- Sandberg, A., and Bostrom, N. (2008). *Whole Brain Emulation: A Roadmap* (No. 2008-3). Oxford University, Future of Humanity Institute.
- Schatsky, D., Muraskin, C., and Gurumurthy, R. (2014). *Demystifying Artificial Intelligence*. Deloitte University Press. November 4, 2014.
- Schwab, K., and Samans, R. (2016a). The Fourth Industrial Revolution: What it Means, How to Respond. *World Economic Forum*. January 14, 2016. Available: <https://www.weforum.org/agenda/2016/01/the-fourth-industrial-revolution-what-it-means-and-how-to-respond/> (accessed on July 20, 2016).
- Schwab, K., and Samans, R. (2016b). The Future of Jobs: Employment, Skill and Human Capital. *World Economic Forum*. January 2016. Available: https://www3.weforum.org/docs/WEF_Future_of_Jobs.pdf (accessed on May 13, 2016).
- Shankelman, J. (2016). Germany Just Got Almost All of Its Power From Renewable Energy. *Bloomberg*. May 16, 2016. Available: <https://www.bloomberg.com/news/articles/2016-05-16/germany-just-got-almost-all-of-its-power-from-renewable-energy> (accessed on April 28, 2017).
- Simchi-Levit, D. (2012). 2012 Annual Re-shoring Report. Massachusetts Institute of Technology. Available: http://supplychain.mit.edu/wp-content/uploads/2014/10/mit_forum_2012_annual_u_s_resourcing_report.pdf (accessed on April 19, 2017).
- Simonite, T. (2016). Moore's Law is Dead. Now What? *MIT Technology Review*. May 13, 2016. Available: <https://www.technologyreview.com/s/601441/moores-law-is-dead-now-what/> (accessed on February 4, 2017).
- Sirkin, H., Zinser, M., and Rose, J.R. (2015). The Robotics Revolution. The Boston Consulting Group, Boston, Massachusetts. Available: http://library.kiet.re.kr/_MultiData/krms/00040/057/The%20Robotics%20Revolution.pdf (accessed on July 12, 2016).
- Stone P., Brooks, R., Brynjolfsson, E., Calo, R., Etzioni, O., Hager, G., et al. (2016). Artificial Intelligence and Life in 2030. One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel. Stanford University, Palo Alto, California. Available: <http://ai100.stanford.edu/2016-report>.
- Stewart, I., De, D., and Cole, A. (2015). Technology and People: The Great Job-Creating Machine. Deloitte. Available: <http://www2.deloitte.com/uk/en/pages/finance/articles/technology-and-people.html> (accessed on May 17, 2016).
- Stiroh, K.J. (2002). Information Technology and the US Productivity Revival: What do the Industry Data Say? *American Economic Review*, 92, 1559-1576.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. (2013). Intriguing Properties of Neural Networks. in *Computer Vision and Pattern Recognition*, 2013.
- Taigman, Y., Yang, M., Ranzato, M.A., and Wolf, L. (2014). DeepFace: Closing the Gap to Human-Level Performance in Face Verification. In *Conference on Computer Vision and Pattern Recognition*. June 2014. Columbus, Ohio.
- The Economist*. (2013). Coming Home. *The Economist*. January 19, 2013. Available: <http://www.economist.com/news/special-report/21569570-growing-number-american-companies-are-moving-their-manufacturing-back-united> (accessed on July 12, 2016).
- The Economist*. (2015). The Sky's the Limit. *The Economist*. October 17, 2015. Available <http://www.economist.com/news/leaders/21674714-shifting-computer-power-cloud-brings-many-benefitsbut-dont-ignore-risks-skys-limit> (accessed on April 2, 2017).
- The Economist*. (2016). Why Firms are Piling into Artificial Intelligence. *The Economist*. April 1, 2016. Available: <http://www.economist.com/blogs/economist-explains/2016/04/economist-explains> (accessed on March 3, 2017).
- The Economist*. (2016a). After Moore's Law. *The Economist*. March 12, 2016. Available: <http://www.economist.com/technology-quarterly/2016-03-12/after-moores-law> (accessed on February 28, 2017).
- The Economist*. (2016b). The Trouble with GDP. *The Economist*. April 30, 2016. Available: <http://www.economist.com/news/briefing/21697845-gross-domestic-product-gdp-increasingly-poor-measure-prosperity-it-not-even> (accessed on April 28, 2017).

- The Economist*. (2017). Finding a Voice. *The Economist*. Available: <http://www.economist.com/technology-quarterly/2017-05-01/language> (accessed on February 21, 2017).
- Thompson, A. (2016). Germany Is Using AI to Smooth the Fluctuations in Its Power Grid. *Popular Mechanics*. Available: <http://www.popularmechanics.com/technology/infrastructure/a21826/germany-machine-learning-power-grid/> (accessed on April 19, 2017).
- Thompson, S. (2016). GDP a Poor Measure of Progress, say Davos Economists. *World Economic Forum*. January 23, 2016. Available: <https://www.weforum.org/agenda/2016/01/gdp/> (accessed on April 28, 2017).
- Tractica. (2016). Computer Vision Hardware and Software Market to Reach \$48.6 Billion by 2022. *Tractica Research*. 2 June 20, 2016. Available: <https://www.businesswire.com/news/home/20160620005440/en/Computer-Vision-Hardware-and-Software-Market-to-Reach-48.6-Billion-by-2022-According-to-Tractica> (accessed on April 28, 2017).
- Turing, A.M. (1950). Computing Machinery and Intelligence. *Mind*, 49, 433-460.
- UNCTAD. (2015). Key Statistics and Trends in International Trade. UN Conference on Trade and Development, Geneva, Switzerland.
- Vere, S.A. (1992). A Cognitive Process Shell. *Behavioral Brain Science*, 15, 460-461.
- Vincent, J. (2016). This Warehouse Robot Just Won Amazon's Shelf-Stocking Challenge. *Popularmechnics*. July 5, 2016. Available: <https://www.theverge.com/2016/7/5/12095788/amazon-picking-robot-challenge-2016> (accessed on May 1, 2017).
- Waters, R. (2015). Investor Rush to Artificial Intelligence is the Real Deal. *Financial Times*. January 4, 2015. Available: <http://www.ft.com/intl/cms/s/2/019b3702-92a2-11e4-a1fd-00144feabdc0.html#axzz49PRkdWgc> (accessed on May 12, 2016).
- Waters, R., and Bradshaw, T. (2016). Robot Economy Sparks Global Investment Boom. *Financial Times*. May 3, 2016. Available: <https://www.ft.com/content/5a352264-0e26-11e6-ad80-67655613c2d6> (accessed on May 6, 2017).
- Wilson, G., and Staffell, I. (2017). The Year Coal Collapsed: 2016 was a Turning Point for Britain's Electricity. *The Conversation*. January 6, 2017. Available: <https://theconversation.com/the-year-coal-collapsed-2016-was-a-turning-point-for-britains-electricity-70877> (accessed on May 1, 2017).
- World Bank Group. (Ed.). (2016). Digital Dividends, World Development Report. International Bank for Reconstruction and Development / The World Bank, Washington, DC.
- Wong, J. (2016). The Workers in these Countries Believe AI and Robots will Replace Them. *Quartz*. March 18, 2016. Available: <https://qz.com/642741/the-workers-in-these-countries-believe-ai-and-robots-will-replace-them/> (accessed on March 23, 2017).
- Xiong, W., Droppo, J., Huang, X., Seide, F., Seltzer, M., Stolcke, A., et al. (2016). Achieving Human Parity in Conversational Speech Recognition. Microsoft Research (Technical Report MSR- TR-2016-71). Revised February 2017. Available: <https://arxiv.org/abs/1610.05256> (accessed on March 23, 2017).
- Yaraghi, N., and Ravi, S. (2016). The Current and Future State of the Sharing Economy. Brookings Institute. December 29, 2016. Washington, DC.
- Zion Market Research. (2017). Global Service Robotics Market Expected to reach USD 24.10 Billion by 2022. Zion Market Research. February 1, 2017. Available: <https://globenewswire.com/news-release/2017/02/01/912700/0/en/Global-Service-Robotics-Market-expected-to-reach-USD-24-10-Billion-by-2022-Zion-Market-Research.html> (accessed on April 28, 2017).