



EPIC



RESEARCH REPORT

Prepared by EPIC Research & Development

A 'STEEP' ANALYSIS REGARDING THE STATE OF ARTIFICIAL INTELLIGENCE 2018/2019

Prepared by Dale P. Deacon
Epic Research & Development



Table of Contents:

Introduction.....	4
Terminology.....	5
A Brief History of Artificial Intelligence.....	6
Understanding the issue from multiple perspectives.....	7
A Systems Perspective.....	7
The STEEP Framework.....	7
Part 1: Environmental Factors.....	8
Environmental: A Perspective on Global Challenges.....	8
Environmental: The Sustainability Challenge.....	9
Part 2: Economic Factors.....	10
Economic: Positive Macro Economic Trends.....	10
Economic: The Western Intelligentsia.....	11
Economic: Global Leadership.....	12
Economic: The Rising East.....	13
Economic: New Businesses.....	14
Part 3: Socio-cultural Factors.....	15
Societal: The Future of Work & Happiness.....	15
Societal: Gender Gap.....	16
Societal: Sentiment Toward AI.....	17
Part 4: Political Factors.....	18
Political: Policy Formation.....	18
Political: Comparing China and the US.....	19
Part 5: Technological Factors.....	20
Technological: Philosophy.....	20
Technological: Moore’s Law.....	21
Technological: Information asymmetry.....	22
Technological: Generalized Artificial Intelligence.....	23
Conclusion & Recommendations.....	24
Bibliography.....	25

INTRODUCTION

It is the year 2035. Around the world, one of the pressing issues of the day, is that of artificial intelligence. Specifically, people's relationship with it, and hence, with the very nature of intelligence itself. The question many ask is,

How should human beings and their machines co-habitate now, and in the near future?

The once-clear lines between Homo sapiens and their creations have blurred significantly. In 2019, cellphones allow us to tap into a global knowledge base, that is Wikipedia, at a moment's notice. Wearable devices monitor our hearts, our breathing, and our activity levels to aid us in living longer, healthier, happier lives. And for some of us, machines are essential to life itself, as is the case with the millions of people, from all walks of life, that rely on pacemakers to survive.

Technologies ubiquitously pervade our world and take a myriad of shapes and sizes. Some machines are ephemeral software programmes that guide machinery. Some take the form of minute sensors and silently gather data. Others still, are more biologically engrained in us, such as hearing aids and artificial hip joints. This litany of devices are loosely categorized under the single term, *technologies*. We'll discuss mechanical technologies in this report, but the main focus will be on information technologies, an subset of which, is that of self-improving information systems. This group of technologies is sometimes referred to as *artificial intelligence (AI)*. What exactly does this mean? To get a clearer idea of this concept, we can turn to an iconic piece of work in this field. In *Artificial Intelligence, A Modern Approach* written by Stuart Russell and Peter Norvig, the duo share a number of definitions uncovered in the literature. After considering a number of popular definitions that have been used over the years, they themselves put it that:

“[AI] attempts not just to *understand* but also to *build* intelligent entities.”

This report will focus on understanding intelligent agents, specifically those that are man-made, and their relation to their makers. Since information technologies often seamlessly, and without restriction, cross international boundaries, we'll be analysing the topic from a global perspective, rather than focusing on localized effects and implications.

Terminology

Taking a step back, we can appreciate that technological fields have a reputation for seeming rather intimidating to those on the periphery, or outside the field. This may be in part due to a swathe of technical jargon. As alluded to on the previous page, this is might be due to the fact that we tend to categorize all technology-related fields into a single category.

Within the broad category of “technology” there are hundreds of thousands of sub-fields, each with their own technical language. This same problem is exacerbated when we narrow our gaze to the technological sub-fields that comprise the term “artificial intelligence”.

As depicted in Figure 1, Artificial Intelligence currently admits a large variety of sub-fields, from general-purpose intelligence, to specific “narrow” fields such as playing games, creating novel artworks, or navigating a vehicle. (Russell & Norvig, 1995/2009)

Artificial intelligence is an umbrella term for a number of computer science technologies, specifically those that attempt to emulate and improve upon tasks regularly performed by the human brain. These systems make use of sensors, cameras and other input to gather information. These systems then rely on huge sets of collected data to continuously improve its own capabilities. (Statista, 2018)

Why is this worth investigating, though? AI is a relatively new field within the sciences and engineering, which means that there are many opportunities to make original findings and contribute towards a growing industry. Similarly, artificial intelligence rest upon some philosophical underpinnings of what is is to perceive, understand and act. These concepts are influenced by some of the great philosophical thinkers throughout history. (Russell & Norvig, 1995/2009) This work is liable to be interpreted in a myriad ways, and the interpretations of which are well suited to a philosophically-minded inquiry that analyses the systemic implications on many other aspects of life, the economy and culture.

To better understand where this technology will lead us in 2035, we will need to (briefly) explore the historical lineage of the field in this report. After which, we will analyse the technological, societal, environmental, and economic implications of these phenomena.

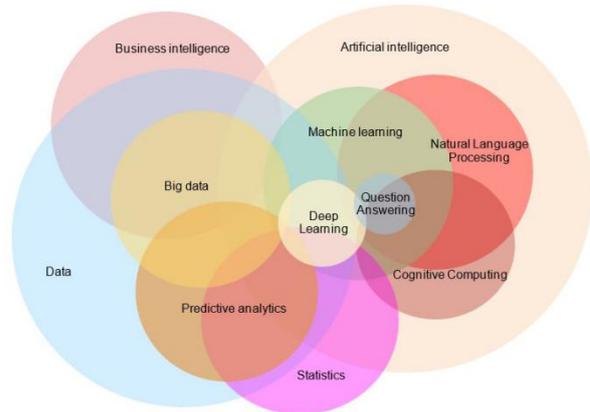


Figure 1: The many sub-fields and overlapping parallel fields of data science and artificial intelligence.

A Brief History of Artificial Intelligence

1943 Gestation of AI: The first work that has been recognized as Artificial Intelligence was performed by Warren McCulloch and Walter Pitts. They drew upon three sources: basic physiology, propositional logic, and computation.

1956 “AI” Coined: John McCarthy coined the term “artificial intelligence”, after convincing US researchers in automata theory, neural networks, and the study of intelligence to coordinate efforts.

1980 Industrial AI: The first successful commercial AI system, known as R1, began operation at the Digital Equipment Corporation. The program helped configure orders for new computer systems. (Russell & Norvig, 1995/2009)

1997 Chess: In 1997 IBM’s DeepBlue system beat grand master chess champion Gary Kasparov. Today, similarly skillful programs that are able to run on smart phones can play at a similar level.

2011 Jeopardy!: IBM’s Watson system competed in the general knowledge quiz show, Jeopardy! against two former winners of the show. The IBM system won the competition by a large margin and claimed the first place prize of \$1 million.

2015 Atari Games: In 2015, DeepMind, a UK based company owned by Google, used a reinforcement learning technique to help a machine learn how to play over forty Atari video games, simply by observing the pixels on the screen. The system was then able to surpass human-level performance in a majority of these games.

2016 ImageNet: In 2016, the image identification and labeling algorithm, ImageNet achieved a super-human level of competency. The error rate of automatic labeling of was reduced to less than 3%, where the average human performance remained at roughly 5%.

2016 Go: In March of 2016, the Google DeepMind team made history again. Their AlphaGo system beat one of the world’s greatest Go players, Lee Sedol. DeepMind continued to improve the algorithm until the latest iteration, AlphaGo Zero, beat the original AlphaGo system one hundred games in a row.

2018 Translation: A Microsoft machine translation algorithm that achieved human-level quality and accuracy translating news stories from Chinese to English. The test was performed on “newstest2017”, a data set commonly used in automated translation competitions.

2018 Collaboration: In 2018, the organization, OpenAI, created a team of neural networks that defeated amateur human teams at the real time strategy game, Dota 2. The systems was trained by playing centuries worth of games against itself, learning through what is called “self-play”.

(AI Index 2018 (c) 2018 by Stanford University)

The tasks above are narrowly defined and involve AI performance at a human or superhuman level in game playing and other complex tasks. These achievements, while impressive, say little about any capability to generalize artificial cognitive ability to other tasks. We’ll return to this topic later.

UNDERSTANDING THE ISSUE FROM MULTIPLE PERSPECTIVES

A Systems Perspective

From a systems-theoretical perspective, we should acknowledge that all of the people, organizations and nations affected by this technology are in reality, intimately interconnected. Relationships exist, either directly or indirectly, between themselves and the larger environment.

In effect, the internet allows for a relatively free flow of information (and culture) around the globe, whilst globalization continues to integrate national economies of trade on an unprecedented scale. Despite any attempts at compartmentalizing individual aspects of any environment, it is important to remember that an issue is never isolated. Each piece of the puzzle affects, and is affected, by its environment. The full dynamics and complexity within the macro environment are beyond our level of abstract consideration. The outcome of any scan of the entire environment, however, does not necessarily take every incidental factor into consideration, therefore, we must admit that any study along these lines is subject to revision and expansion as new information comes to light. (Van Vuuren, 2016)

The STEEP Framework

That said, it's important that we formulate our ideas about the topic in a systematic manner. To better understand the issue itself, we need to analyse the contextual environment in which it resides.

The environment in which we all operate can be categorized and analyzed using the popular modern adaptation known as the STEEP/PESTLE framework. Frameworks of this kind have been in use for a while, as far back as Francis Aguilar's ETPS framework in 1967. (Choo, 2001)

The STEEP framework entails, at a very broad level, focusing our attention on the Social, Technological, Economic, Environmental and Political categories. From here, we can analyse the implications and potential knock-on effects our core research question may have. Throughout this report, we'll utilize this framework as the basis of our analysis of the topic of artificial intelligence, but will do so in a slightly different order to best build one concept upon another.

Part 1: Environmental Factors

Environmental: A Perspective on Global Challenges

In 2018 Google CEO, Sundar Pichai, said in a Google blog piece,

“At its heart, AI is computer programming that learns and adapts.” He continues, “It can’t solve every problem, but its potential to improve our lives is profound. At Google, we use AI to make products more useful—from email that’s spam-free and easier to compose, to a digital assistant you can speak to naturally, to photos that pop the fun stuff out for you to enjoy.” (Google, 2018)

Despite Pichai’s optimistic disposition, today, in 2019, there are a number of existential threats that the human race is confronted with. Fortunately, there are also a number of institutions dedicated to uncovering, understanding and monitoring these threats.

Among them, the Bulletin of Atomic Scientists, Future of Life Institute, and Future of Humanity Institute feature prominently. Each of these institutions has shared a number of critical global-scale challenges that the human race needs to address. The 2018 Doomsday Clock Statement stated that “It is now two minutes to midnight - the closest the Clock has ever been to Doomsday, and as close as it was in 1953, at the height of the Cold War.” (Bulletin of the Atomic Scientists, 2018)



Figure 2. The Doomsday Clock at “Two minutes to midnight”

(Bulletin of the Atomic Scientists, 2018)

The Bulletin discusses a number of existential threats that we will need to navigate. Their list describes the following:

- I. The uncoordinated response to *climate change*; According to the bulletin, last year, the US government pursued ineffectual policies on climate change. This in stark contrast with prior commitments to curb carbon emissions and modernize the country’s energy policy.
- II. The *nuclear threat*; the risk that nuclear weapons may be used either intentionally or by miscalculation. According to the bulletin, this threat grew last year around the globe due to political tensions in North Korea, Pakistan, India, and The United States.
- III. And *emerging technologies*; The concern over a lack of trust in political establishments, in online and traditional media, in scientific expertise. This has been, in part, attributed to an abuse of information technologies.

(Bulletin of the Atomic Scientists, 2018)

We will discuss the first among these issues in this section as a global environmental factor. Across the majority of this report, we will assess and discuss emerging technologies.

Environmental: The Sustainability Challenge

Importantly, environmental collapse, as described by the Club of Rome's Limits to Growth report of 1972, needs to be avoided. This is no small task. We will have needed to address shortages in the areas of food production, arable land, clean air, clean water, non-renewable resource consumption, and others. Each of these variables is comprised of a litany of complex sub-challenges. (Meadows et. al, 1972) This may sound overwhelming, but with a global network of creative thinkers at our disposal, the population of Earth is rising to the occasion.

For example, the overpopulation trend that was so pointed in the Limits to Growth report may not be as dire as we once believed, thanks to insight shared by Hans Rosling among others. "In 1948, women on average gave birth to *five* children each. After 1965 the number started dropping like it never had done before. Over the last 50 years it dropped all the way to the amazingly low world average of just below 2.5." [Emphasis added.] (Hans Rosling, Anna Rosling Rönnlund, Ola Rosling, 2018)

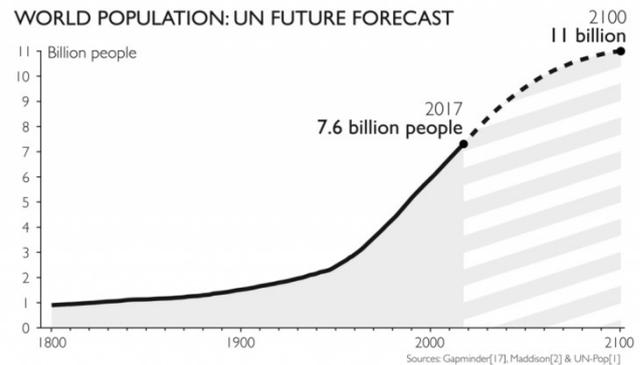


Figure 3. The leveling off of the global population as according to Rosling's Gapminder Foundation.

<https://www.gapminder.org/topics/population-forecasts/>

The dramatic decrease in the number of children per woman is expected to continue. This should remain true so long as more people continue to make progress in the areas of poverty alleviation, education, contraceptives, and sexual education. (Hans Rosling, Anna Rosling Rönnlund, Ola Rosling, 2018) Additionally, a sustainable energy policy, that drastically reduces our reliance on hydrocarbon fuels, will have had to have been created and implemented. Professor Klaus Schwab, Founder and Executive Chairman of the World Economic Forum, wrote in his book, *The Fourth Industrial Revolution* that a "Circular Economy" more directed toward the reuse and maintenance of consumer goods is necessary. This, in opposition to the idea that mass-production and consumption.

"At the heart of this promise is the opportunity to shift businesses and consumers away from the linear take-make-dispose model of resource use, which relies on large quantities of easily accessible resources, and towards a new industrial model where effective flows of materials, energy, labour and now information interact with each other and promote by design a restorative, regenerative and more productive economic system." (Klaus Schwab, 2016)

These challenges are immense, but may potentially be overcome. Trend extrapolation and simulation models rarely consider the impact of unprecedented, radical, and coordinated intervention by numerous businesses and states.

Part 2: Economic Factors

Economic: Positive Macro Economic Trends

The Gross Domestic Product (GDP) of an economy is a standardized measure of total production within a state. More precisely, it is the total value of all goods and services produced.

Growth is then measured by the rate of change in GDP from one year to the next. Through a long-term perspective, it's clear that prosperity and consistent growth are recent achievements.

(Max Roser, Our World In Data, 2019)

This sheer increase appears to take the shape of unprecedented growth. In the 1972 Limits to Growth report, The Club of Rome team explained that a quantity can be said to exhibit exponential growth where it “increases by a constant percentage of the whole in a constant time period.” (Meadows et. al. 1972)

Some technologists, such as Ray Kurzweil, consider this steep curve to be a by-product of technological advancement. Kurzweil has argued that technological advancement is analogous to evolutionary processes. He points to how the evolution of biological species gave rise to the first technology-creating species, human beings. This then, paved the way for a new evolutionary process of directed, technological evolution. The apparently exponential increases in GDP and productivity could be attributed to what has been in Kurzweil's work referred to as the *law of accelerating returns*. (Kurzweil, 2001) This should result in greatly increased productivity, then.

Per capita measurements divide the GDP output by a country or region's population. This allows us to evaluate levels of productivity over a certain period of time. Using this measure, we could similarly observe an average level of productivity while discounting outliers among the population. However, labour productivity seems to be closely analogous with capital-intensive financial activity that has potentially large rewards in exchange for relatively small time investment. This makes measuring labour productivity in real terms difficult.

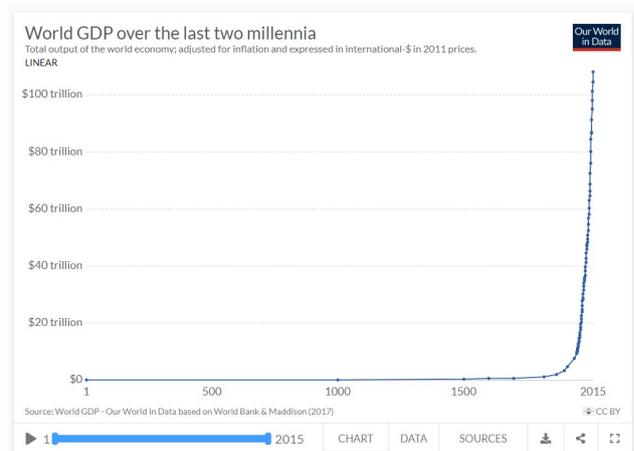


Figure 4. Global GDP over a 2000 year time-frame. (Max Roser, Our World In Data, 2019)

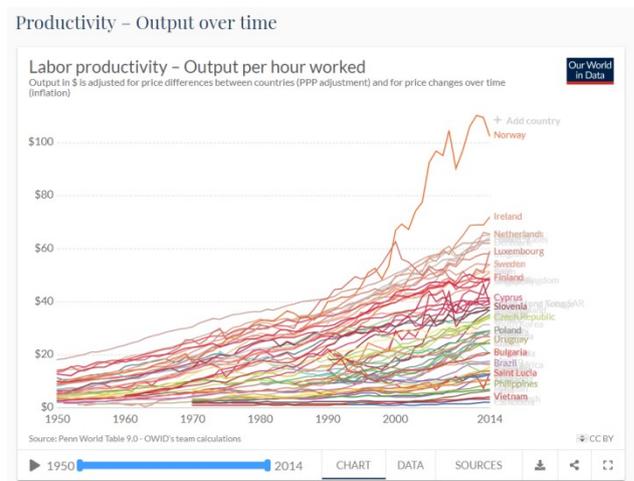


Figure 5. Labour productivity over time since 1950. (Max Roser, Our World In Data, 2019)

Economic: The Western Intelligentsia

As of June 2018, there were 4,925 artificial intelligence enterprises worldwide. The establishment of new artificial intelligence companies has slowed worldwide since 2015 when the number of newly established enterprises across the globe reached 847. (Statista, 2018)

Much of Western Europe and Scandinavia share a strong economic positions with their developed, services-based economies. As such, the UK, Germany, and France host a burgeoning AI industry. (Statista, 2018)

As illustrated in Figure 6, European firms and governments however, lag behind in size and scale when compared to the American AI industry. Decades of investment in information technology within the US has nurtured a thriving technologically-enabled service industry.

Number of AI companies worldwide 2018, by country

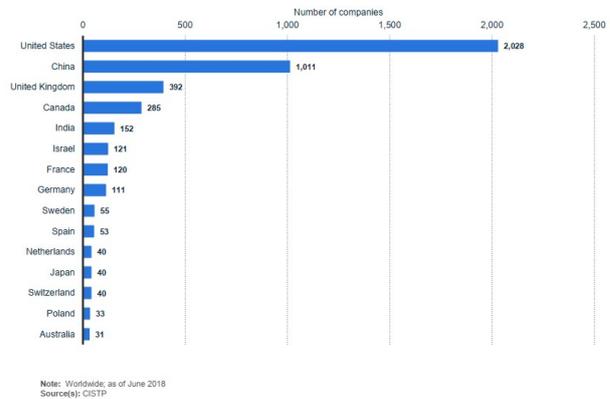


Figure 6. Number of AI companies worldwide as of June 2018, by country (Statista, 2018)

Despite some periodic economic fluctuations, American firms have garnered immense growth over the past fifty years. This has endowed American businesses, and by extension its government, with an abundance of financial capital, intellectual property, and skilled labour.

This has allowed the US to lead pioneering initiatives in the information technology, and artificial intelligence industries for decades. As figure 7 shows, the US remains in a dominant economic position to this day with the majority of AI businesses finding their homes here. (Statista, 2018)

Share of AI companies as of June 2018



Figure 7. The share of AI businesses worldwide as of June 2018 (Statista, 2018)

Within the American technology industry, there are a number of major players in the artificial intelligence field. Predominant among them, a handful of the highest valued companies in the world. Trade giant, Amazon, has invested heavily in both the consumer and business oriented AI. Their cloud service, Amazon Web Services (AWS) makes available a number of machine learning programs and pre-trained AI services for businesses. The company currently has over 10,000 diverse customers, including Netflix, NFL, and NASA. Alexa, Amazon’s won AI language assistant, is integrated into its consumer electronics. Microsoft too, is involved in Artificial Intelligence and caters to the needs of both consumers and businesses. Cortana, Microsoft’s artificially intelligent digital assistant, competes alongside Amazon’s Alexa, Apple’s Siri, and the Google Assistant. As is the case with Amazon, artificial intelligence tools, such as chat bots and machine learning algorithms, play a major role in Microsoft’s Azure Cloud service.

Economic: Global Leadership

Perhaps the largest AI company at the moment is search giant, Google. In a public Google blog post published in 2018, CEO Sundar Pichai shared AI principles upon which Google operate. “Google invests heavily in AI research and development, and makes AI technologies widely available to others via our tools and open-source code.” (Google, 2018)

The capabilities of Google’s AI research and development has continued to grow as they acquire AI-focused start-ups as well as potential competitors. Over recent years, the company has created or purchased no fewer than twelve new artificial intelligence businesses. Importantly, this includes a \$400 million deal to acquire DeepMind, and hence, the team that formulated the Go-playing algorithm, AlphaGo.

Number of artificial intelligence startups acquired from 2010 to June 2018, by company
AI startup acquisitions by company 2010-2018

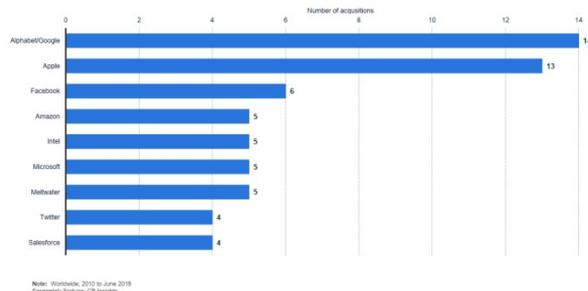


Figure 8. Artificial intelligence acquisitions from 2010 to June 2018. (Statista, 2018)

As far as laying the groundwork for machine learning software development goes, Google appears to lead the field.

Popularity/usage of artificial intelligence frameworks worldwide 2018

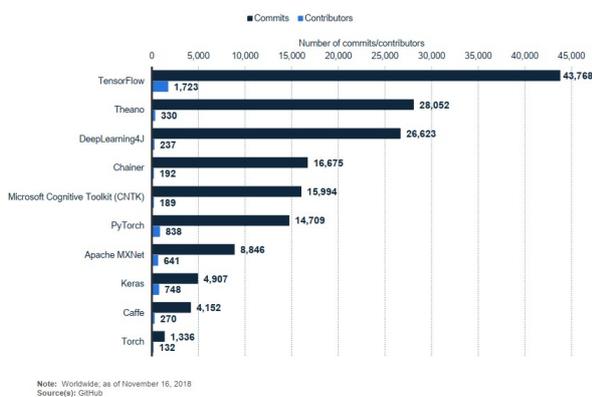


Figure 9. Google’s TensorFlow machine learning platform currently leads in the number of developers and projects according to GitHub. (Statista, 2018)

Google’s machine learning development platform, TensorFlow, is more popular and widely used than any other in the same category. This may be in part due to the fact that it is currently free to download and learn.

What this means is that many developers around the world are learning to use tools specifically tailored towards growing Google’s influence within this growing industry. The unpaid learning of Google’s tools may result in a large proportion of skilled developers at Google’s disposal should they require them in the near future.

Economic: The Rising East

At the other end of the world, a globalized economy has generated much growth and prosperity for the majority of the world's population living in India and China. With a bustling industrial sector and rising standards of living, China is setting the stage to become a technological superpower as well. In July 2017, the State Council of China released the Next Generation Artificial Intelligence Development Plan which outlines the nation's approach for AI up to the year 2030. (Statista, 2018)

The scale with which China has invested in AI research and development is in the process of significantly changing the balance of influence in this field. With a burgeoning AI industry, China aims to provide a viable alternative to US engineered solutions in this space.

Share of global AI investment and financing from 2013 to Q1 '18

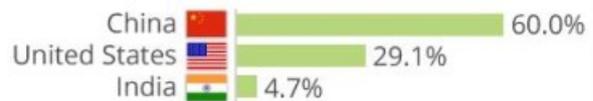


Figure 10. Share of global AI investment and financing from 2013 to 2018. (Statista, 2018)

The Next Generation Artificial Intelligence Development Plan involves keeping pace with the latest developments in artificial intelligence while closing the gap between itself and the United States.

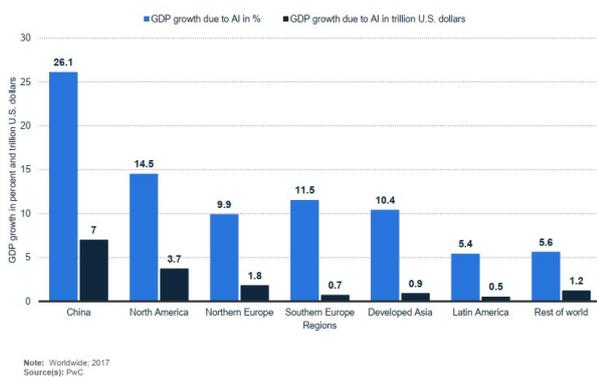


Figure 11. Increase of GDPs globally due to artificial intelligence by the year 2030

As illustrated in figure 11, consulting firm, PwC estimates that roughly 26,1% of all GDP gains to be had by AI by 2030 will have been made in China. This, in stark contrast with the next biggest beneficiary, North America at 14,5%.

According to Statista, Chinese companies, Baidu, Alibaba, Tencent and iFlytek have all been recognized by the Chinese government as leading figures in the ongoing development of AI technologies. (Statista, 2019)

Highest valued artificial intelligence companies in China in 2018 (in billion U.S. dollars)
Artificial intelligence unicorn companies in China 2018, by value

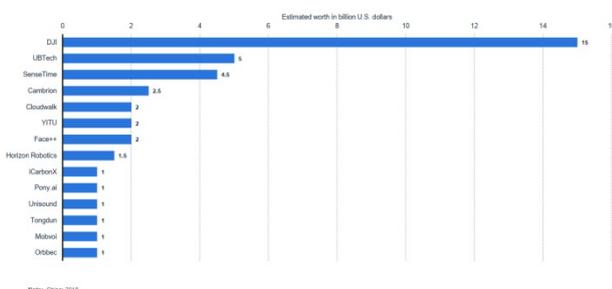


Figure 12. Chinese company DJI features as the highest valued artificial intelligence company in China in 2018.

Chinese companies are already taking significant strides ahead in AI. The drone manufacturer, DJI, has already been valued at \$15 billion and the company has a reported share of over 70% of the global drone market. DJI is increasingly pursuing an interest in the AI industry, as their drones use AI-based image recognition to avoid colliding with nearby objects. With significant investment and industry leading technologists, Chinese prospects for leading the global AI industry by 2035 appear strong. (Statista, 2019)

Economic: New Businesses

New technologies create opportunities for entrepreneurs, provided they are able to receive sufficient funding and support.

These entrepreneurs are able to formulate new models and products that disrupt the balance established industries.

The AI industry is no different. The number of new businesses in the artificial intelligence industry has boomed in the last two decades. Figure 13 shows the number of active venture-backed US private start ups between 1995 and 2018. The solid blue line represents AI start ups only, while the grey line shows all venture-backed start ups, including AI start ups. (AI Index, 2018)

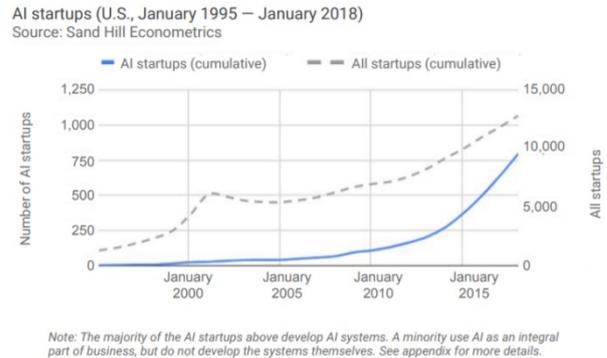


Figure 13. Active AI start ups in the U.S. increased 113% from 2015 to 2018 while active start ups as a whole increased 28%. (AI Index, 2018)

For the most part, growth in all active start ups has remained relatively steady throughout this period. The number of AI start ups alone however, has seen exponential growth.

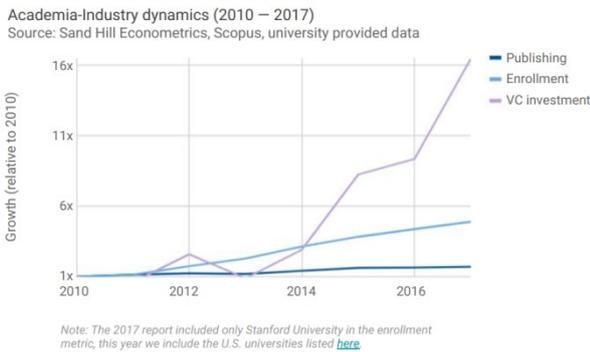


Figure 14. Venture Capital investment within the US has increased over 16 times its 2010 levels. (AI Index, 2018)

Similarly, venture capital funding in this space has increased. In part, due to the renewed hype around the subject.

The “AI winter” has given way to what is being called the “AI spring”. (Statista, 2018) In 2017, AI investment and financing reached roughly \$40 billion dollars globally; A \$10 billion increase on the previous year. (AI Index, 2018)

The growth of investment in AI over the past few years has led to more acquisitions of start ups as well. US technology company, Intel, is one of the most active investors in AI with over one billion USD invested in AI start ups. With an influx of capital, we can expect to see more companies, more leaders and more breakthroughs being made in this space over the next decade. As we’ll see however, with significant economic shifts, there are bound to be many complex societal implications to these decisions.

Part 3: Socio-cultural Factors

Artificial intelligence is already having a profound effect on the share of investment, and hence, employment. As such, there are profound knock-on effects that can be measured from a sociocultural perspective.

Societal: The Future of Work & Happiness

The future of work is a topic of ongoing scrutiny, as these developments impact both economic and social aspects of the environment. The ongoing technology-reliant trend of automation is a major factor in this debate. Automation is naturally considered to be a critical factor driving change in the future job market. As has been evidenced by similar technological shifts in the past, the adoption of automation has the potential to drastically transform the workplace once again. (Statista, 2018)

The rate of change in employment has been relatively constant over the past 50 years. Routine jobs have consistently lost employees to the service industry, which favors non-routine, cognitive tasks.

If machine learning and the automation of cognitive capabilities advance quickly, we could see a drastic increase in productivity in select areas of the economy. This could increase the pace of adoption, and hence, further incentivise further development of cognitive automation functions. (EY, 2018)

As we can see in figure 15, the number of weekly hours worked has been on a steady decline since 1870. (Max Roser, Our World In Data, 2019)

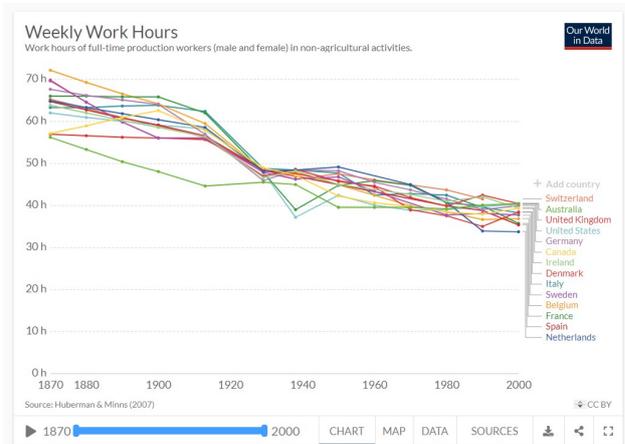


Figure 15. Michael Huberman and Chris Minns published estimates of weekly work hours going back to the late 19th century. This data shows that over this time working hours have steeply declined.

(Max Roser, Our World In Data, 2019)

Declining working hours may very well continue and accelerate along this trajectory. The pace with which this decline proceeds may increase if monetary and political incentives remain as they are. Currently, some of the most accomplished programmers and strategists are on a path that rewards the automation of labour in an attempt to garner greater market share for their employers. Success for these teams would come at the expense of (potentially) millions of working class people's livelihoods. (Executive Office of the President (US), 2016)

Societal: Gender Gap

Not all jobs are lost, however. Many new fields are emerging and growing, and with new businesses, come new employment opportunities. That said, it appears that these new opportunities are not all shared among the population equally. Historically, gender inequality has been pronounced in the technology industry. As we'll see, during this new wave of employment in the artificial intelligence sector, these same historical biases seem to remain.

In the workplace, these figures remain tilted in favour of male participation. Figure 16 shows the proportion of male and female applicants for AI job openings in 2017. According to this Gartner study, on average, men comprise 71% of the pool for AI jobs in the US.

Since the requirement for expertise in the sub-field of Machine learning is commonly sought after, this category has the highest volume of applicants. (AI Index, 2018)

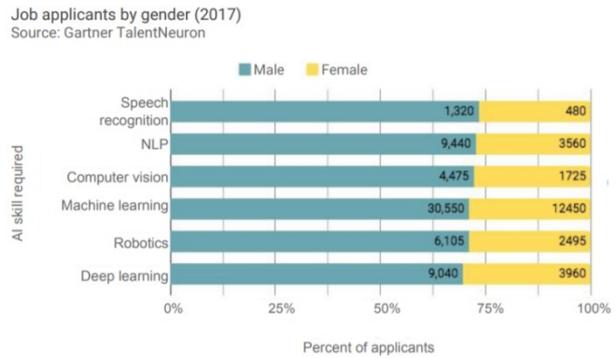


Figure 16. On average, men make up 71% of the applicant pool for AI jobs in the U.S. (AI Index, 2018)

The results found in the workplace were similar to those found in academia. Across the tertiary institutions studied in this report, on average, the vast majority of AI professors were male. There was apparently little variation between schools, regardless of their location.

In the sampled institutions, on average, 80% of professors from UC Berkeley, Stanford, UIUC, CMU, UC London, Oxford, and ETH Zurich are male. (AI Index Report, 2018)

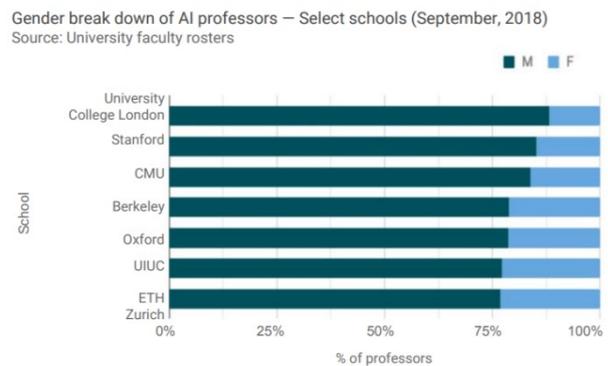


Figure 17. Gender breakdown of AI professors (AI Index Report, 2018)

One issue that could be analyzed further in this regard is the quality and amount of data. Barriers to improved levels of gender diversity in academia often involve a lack of access to data. To this end, institutions have been encouraged to be more transparent. (AI Index Report, 2018) That said, there is little guarantee that more data would disconfirm these findings. Cultural and traditional biases may play a role in the choices young women make regarding their careers. This same issue has been found across all of the STEM fields, and further investigation is required to adequately address the imbalances.

Societal: Sentiment Toward AI

The macro trend of improved standards of living across the world has had a net-positive effect on many people's disposition. (Hans Rosling, Anna Rosling Rönnlund, Ola Rosling, 2018)

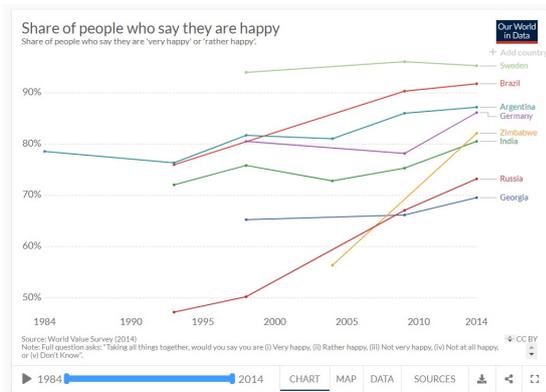


Figure 18. The share of people either Very happy or Rather happy. (Max Roser, Our World In Data, 2019)

As we can see, in figure 18, citizens in the majority of countries have tended to respond positively when asked how satisfied they are with their lives.

In some cases, the improvement has been very great as is the case in Zimbabwe where the share of people who reported being happy grew from 56.4% to 82.1% within a decade. (Max Roser, Our World In Data, 2019)

Likewise, in the media, there has been a significant change in the perception regarding artificial intelligence and machine learning. From Hollywood movies, to publicity stunts, to automated social media bots, artificial intelligences have become a topic of mainstream interest. A growing awareness of this topic seems to be taking shape.

Figure 19 shows media articles containing the term "Artificial Intelligence". The emotive sentiment of each of these articles was then analyzed and the results were categorized as either neutral, positive, or negative pieces. The study showed that a significant percentage of AI-related articles have become much less neutral. Whether by design or not, articles have shown a tendency to be more positive, particularly since 2016. (AI Index Report, 2018)

Sentiment of articles referencing AI (2013–2018)
Source: TrendKite

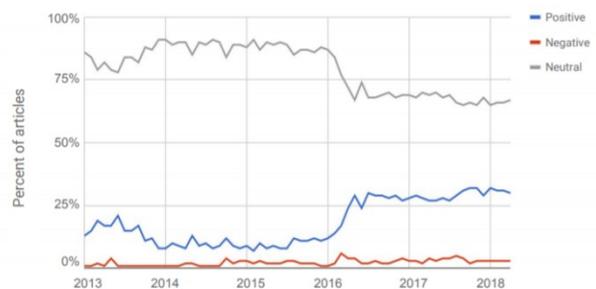


Figure 19. Articles on AI became 1.5x more positive from 2016-2018 (AI Index Report, 2018)

Much of the world's population accesses its news media through portals such as Facebook. The media platform is another of the West's largest and most well-resourced participants in this industry and the company is open about the extent to which they are investigating the application of artificial intelligence. After the events of the "Cambridge Analytica Scandal" became public, Facebook declared that they would focus their artificial intelligence efforts on identifying malicious actors. (Facebook Newsroom, 2017) In effect, the media organization responsible for distributing information to millions of people, had now made a public declaration to rely more heavily on automated AI content management.

Part 4: Political Factors

As society begins to form opinions around these topics, the policies and actions of governments are becoming instrumental in ensuring a smooth adoption process for Industry 4.0 technologies. Many of these technologies have profound social implications, and finding a prudent strategy is fundamental in fostering an environment in which businesses and the public can flourish. Programs that complement these influential digital technologies may help to form a strong support system enabling new companies, and employment opportunities. (EY, 2018)

Political: Policy Formation

Politically, there are many stakeholders involved in understanding, creating, testing, deploying, and monitoring advancements in the artificial intelligence industry. Although we are still in the early days of forming political policies regarding AI, the progress made thus far looks promising. Several nations and independent organizations have already started to develop policies for the governance of AI. Last year, the UAE appointed the world's first Minister of Artificial Intelligence. Similarly, Germany took steps in 2017, to develop the world's first set of regulatory guidelines for automated and connected driving. This progress is noteworthy, yet to date, these efforts have not coalesced into a larger governance framework that is able to extend beyond national boundaries.

The question remains; What type of government involvement would allow for safer and better outcomes?

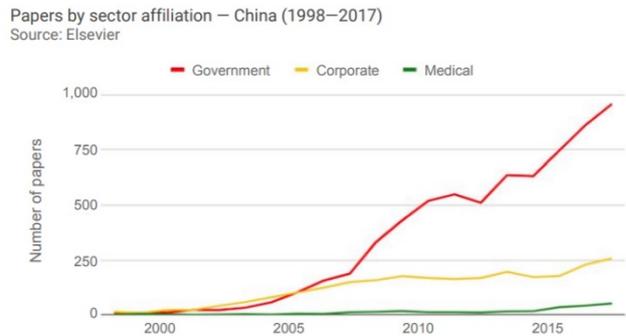
The University of Oxford's Professor Nick Bostrom is the Director of the Future of Humanity Institute. He plays a prominent role in the world's policy formation and believes that in the short-term, governments should be focused on creating formal processes by which new automated technologies can be tested in a safe environment before being brought to market. Bostrom believes that it is important to have a regulatory process for this rather than doing it in a vacuum, to ensure that faults are identified early on before significant damage can be done. In the longer-term, we need a much larger conversation about how we might handle a large worker displacement event as a result of automation. The question of how workers might be supported economically, despite not being able to produce any value to the economy. "The real aim should be to decide, as a society, how to distribute value created by artificial means." (Bostrom & Yudkowsky, 2018)

Bostrom is one of a handful of experts that seek to raise awareness and subsequently resolve these kinds of political issues; By understanding, discussing, and ultimately shaping the policy framework that will ideally guide the development of these technologies in a direction that is beneficial to all, rather than a select few.

Political: Comparing China and the US

Since 2014, the Chinese government has launched a number of national AI initiatives. It has been reported that the goal of the Chinese political establishment is to ensure that China takes a leading role in the global AI industry. Government funding has also been approved for AI-related research projects as well as in a Beijing technology park dedicated to artificial intelligence. (Statista, 2018)

China’s official New Generation AI Development Plan considers a timeline for AI development up until the year 2030. The plan seeks to build China’s AI aptitude to match the most advanced capabilities in the US and Europe by 2020. As figure 20 shows, the Chinese research effort has increased greatly between 2007 and 2017. According to the plan, by 2025, China aims to make major strides in AI policy creation and regulations, once broader use of AI has been established across a variety of industrial sectors. (Statista, 2018)



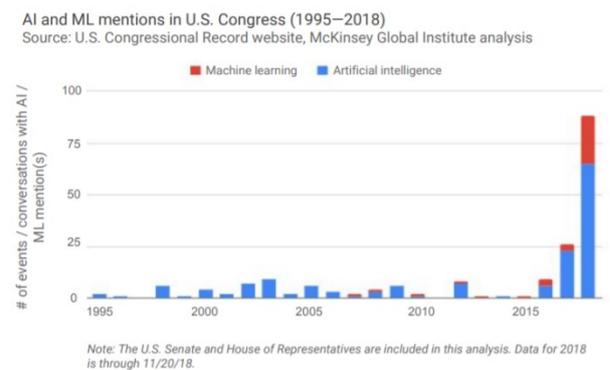
Note: For reference, in 2017 the number of AI papers on Scopus affiliated with the academic sector in China (not shown) was ~14,000 or ~92% of total AI papers in China.

Figure 20. Between 2007 and 2017, government AI research papers in China increased by 400% while corporate AI papers in China increased by 73%. (AI Index, 2018)

Also by 2025, the Chinese plan aims to have made major breakthroughs in AI theory. At this point, AI is expected to become the driving force for industrial upgrading and economic restructuring. And by the year 2030, China plans to become the world’s *primary AI innovation center*. These plans stand in stark contrast with the debate in American congress. While the Chinese one-party system is able to focus efforts on formulating a cogent, long-term strategy, the topic has only recently been publicly brought to the fore in the US. (Statista, 2018)

Figure 21 illustrates how many times the phrases “Artificial intelligence” and “Machine learning” could be found in transcripts of the U.S. Congressional Record and the records of proceedings in the US.

Mentions of both of these terms have increased since 2016. “Machine learning” was rarely mentioned before 2016, and since then has remained relatively rare compared to the term “Artificial intelligence“. (AI Index, 2018)



Note: The U.S. Senate and House of Representatives are included in this analysis. Data for 2018 is through 11/20/18.

Figure 21. AI and ML mentions in U.S. Congress (1995–2018) (AI Index, 2018)

While some official documents outline broad policy ideas, the Executive Office, under the Obama administration, admitted that these developments have the potential to disrupt the ability of millions of Americans to earn a living. (Executive Office of the President (US), 2016)

Part 5: Technological Factors

By the year 2035, the world will have tackled many serious topics that we recognize today. The concerns of 2035 however, may appear quite bizarre and foreign by comparison. MIT professor, Max Tegmark, in his book of the same name, dubbed artificial intelligence “Life 3.0”. (Tegmark, 2017) Tegmark did so because he believed, as Kurzweil did, that technology is some form of the evolutionary process that resulted in tool-wielding primates ushering in their own geological era, or *anthropocene*. When we consider the technological aspects of artificial intelligence between 2019 and 2035, we too need to assess the possibility of whether or not life, and evolution, are truly restricted to organic, biological forms. We will now unpack this complex issue and attempt to use a data-driven approach to separate the important realities of this challenge from those that are trivial or unscientific.

Technological: Philosophy

Artificial Intelligence is a discipline rooted in engineering and computer science, yet its implications are evident in many areas of the economy, society, politics, as well as philosophical ideology. The nature of perception, understanding, cognition, and humanity itself can be explored through the lens of this technology. Ancient Greek philosopher, Aristotle (384–322 B.C.), was the first to formulate laws concerning the rationality of the human mind. He surmised that given some initial premises, we could make use of an informal system of deductive reasoning. This same principle is at work in the logical structuring of computer code that enables automation of informational tasks. Aristotle’s logical framework, in principle, allowed us to generate conclusions mechanically, through step by step progression through a piece of formalized text. Some of these same ideas are used today. Aristotle’s work was in some sense implemented over two thousand years after their conception by Newell and Simon in their GPS system. Today, we refer to it as a “regression planning system”. (Russell & Norvig, 1995/2009)

A common philosophically-inclined concern is that major breakthroughs in software (programming) are made, and then subsequently kept hidden from the rest of the world. This might give a single individual or group an unprecedented amount of influence over the economy, society, politics or a combination of all of them. (Tegmark, 2017) Early AI ethics conferences, such as the Asilomar conference, have led to a widely-respected theory that all benefits of AI should benefit the world’s population, rather than overwhelmingly empower a single entity. To this end, the former non-profit research group, OpenAI, was created. It is oriented around the development of AI for the benefit of all humankind. It has managed to maintain its open source ideology until recently. The project has attracted influential thinkers and investment sponsors alike. This is an example of an informational and technological trend that may greatly influence how the future plays out. If we collectively agree that this kind of incentive is greater than the monetary incentives we currently tend to pursue, a much more amicable implementation of artificial intelligence may be realized.

Technological: Moore's Law

The development of AI depends on the advancement of underlying technologies. Moore's Law, named after, Gordon Moore, is an observation that "the number of transistors on integrated circuits doubles approximately every two years."

This principle seems to provide support for the existence of something approximating a predictable *law of nature* underlying the development of information technologies. Moore's law has been scrutinized and it seems that it similarly applies to technical metrics as diverse as processing speed, product pricing, and memory capacity. They all appear to have been progressing in an exponential manner. Using this observation as a basis, we can consider the implications of this trend if it were to continue into the future. (Max Roser, Our World In Data, 2019)

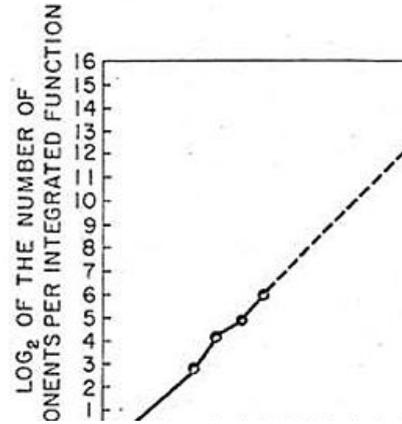


Figure 22. Number of components per integrated function for minimum cost per component extrapolated vs. Time. (Log scale) (Max Roser, Our World In Data, 2019)

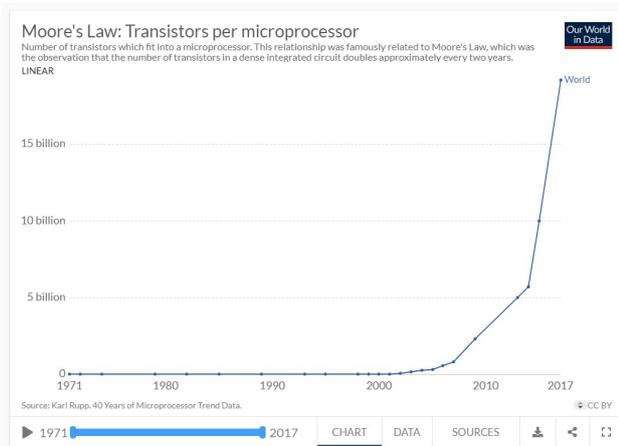


Figure 23. Number of transistors which fit into a microprocessor from 1971 to 2017. (Max Roser, Our World In Data, 2019)

Today, we have reached what appears to be a possible limit to Moore's law. When we work at an atomic scale, we find that minute forces prevent us from decreasing the size of effective semiconductors. Kurzweil does not see this as a problem. He claims that this apparent limitation could be overcome through what he called a "paradigm shift". A fundamental change in the way processing is performed. He describes an overhaul of the process itself, which in effect is what the cutting-edge fields of "parallel processing" and "quantum computing" attempt to achieve. (Kurzweil, 2001)

Should Moore's Law continue, the exponential trends discussed throughout this report may continue. Artificial intelligence that utilizes the processing capabilities of a quantum computer are conceptually profound. However, such an event may not transpire. Kurzweil's exponential curve may simply turn out to be the initial segment of an "s" curve that slowly reaches a limit, or that of a standard bell curve. To better understand the probability of each scenario, further in-depth investigation is required.

Technological: Information asymmetry

Gaining access to some machines is far more challenging than others and cost implications play a major role. Physical transportation and logistics also limits the spread of certain technologies. The list goes on. These factors have resulted in a deeply stratified technological environment.

This same challenge exists in respect to the accessibility of information technologies. There are public institutions that make their research findings available to everyone, but conversely, there are also a number of competing businesses that safeguard their findings. This is to be expected from those that operate through the monetization of information technologies. Companies like Google, Twitter, and Facebook are large enterprises that serve nothing but data to their customers. This means that their products take the form of propriety information. While we transition more and more towards an information economy, we are slowly gaining an appreciation for data as a valuable commodity.

As has been witnessed in the past, one cannot expect intellectual property laws will be acknowledged, respected or enforced, especially across national borders. For this reason, many of the world's most influential and cutting-edge teams secretively guard their operations, and their valuable insights. This means that while useful information in the public domain does exist (for the subset of us that seek it out) there is still vast wealth of information stockpiled by some of the world's most influential companies. To this end, the words of American-Canadian fiction writer and essayist, William Gibson said in *The Economist* in 2003, "The future is already here – it's just not evenly distributed".

This idea is at work in the world of information technologies, the subset of technologies that pertain to the transfer of data. To Gibson's point, it is clear that not all entities are equally well-informed. Nowhere was this made more apparent than during the congressional interrogation of Facebook in 2018. The media organization spent months defending itself against criticism and allegations that it had failed to intervene on issues such as the spread of fake news and extremist propaganda. What struck many was the inability to which many senators were unable to understand or articulate the concerns of their constituents. Some entities simply possess more information than others. As such, those with less information are not aware of important issues to the same extent as well-informed entities. As a result of this *informational asymmetry*, not all entities have equal influence over how the future will transpire. Plans are often made with imperfect information, and therefore need to be malleable in order to accommodate to a rapidly changing informational landscape. A more proactive approach may be to build this concept into the very framework of our systems and processes.

Technological: Generalized Artificial Intelligence

Some of the debate around how we should think about AI in relation to ourselves, centers around the hypothetical scenario of what is called “generalized AI”; an artificial intelligence that is able to perceive, reason and act in a manner that resembles that of a human being. Today in 2019, there are many “narrow AI” systems that are in wide use, but there are no general AIs. These general, or “Strong AI”, systems would exhibit human-levels of intelligence by approximating or mimicking human thinking. (Statista, 2018)

There is almost universal agreement among modern AI professionals that Artificial Intelligence falls far short of human capabilities in one sense or another. This, despite the fact that AI algorithms have surpassed humans in many specific abilities such as playing chess. That said, chess was previously considered to represent the epitome of intelligence until IBM’s DeepBlue beat Kasparov. General intelligence’s defining characteristic is generality, an ability to apply intelligent reasoning to novel circumstances. Current AI algorithms are characterized by a narrowly defined competencies in very restricted domains of expertise. DeepBlue may have become the world chess champion, but it could not play any other game, let alone apply its cognitive powers to a creative or complex task. (Bostrom & Yudkowsky, 2018)

The prospect of AIs with intelligence surpassing our own is not outlandish. Google’s ability to serve information quickly to any user anywhere in the world has been harnessed for some time now. The extraordinary challenge of creating an adaptive artifact that is able to generalize its ability to solve problems is one that yet eludes us. These challenges may seem far-fetched, but when we consider the rapid rate our information systems have advanced over the past fifty years, it seems likely that we will encounter previously outlandish moral and ethical questions either sooner or later. No matter the timing of this issue, it would be prudent to investigate possible scenarios and strategies as soon as we’re able to conceive of them.

CONCLUSION & RECOMMENDATIONS

The challenge we're faced is an existential one. How should we think about ourselves and our creations? We could treat these creations like simple tools, as we have done in the past. This paradigm may come under scrutiny shortly, though. The rate of advancement in this field is staggering. Since many of us are increasingly physically and mentally augmented by our mechanical creations, we could consider treating machines as equals. Today, we may wonder if saying "please" and "thank you" is a moral imperative when interacting with artificial voice assistants. In the near future, far more complex versions of these same systems will be in production. More realistic, and more human-like than ever before. Alternatively, we could regard our creations as our superiors. Machines have far outstripped our collective capabilities in some domains like the games of chess and go. Google's search algorithm is able to fetch relevant information to a query far quicker and more reliably than any human team could. Its superhuman abilities are downplayed since it is restrained to only a single function. Prior to the Google search algorithm, we felt no threat to our usefulness by the desk calculators that were far superior in their ability to calculate mathematical operations. They too could only perform a very limited set of tasks. How might we react when a machine is able to handle a variety of human-like jobs? Machines have the potential to perceive, understand, and act in ways almost unimaginable to the mind of a solitary human being.

As technologists, we may lean towards appreciating machines as far more than mere tools, however, as a systems thinkers, we must realize that machines will be regarded in different ways by different people. A very wide spectrum of reactions will be observed from hate, to love and pride. The Obama Administration painted a number of possible strategies for thinking about this issue: To invest in, and develop AI for its many benefits is essential. As with almost any new technology, there are many risks and opportunities involved. Risks in this case, may take the shape of autonomous attacks or subversive intelligence gathering. Conversely, opportunities lie in using these new cognitive resources in an efficient manner. Instead of having to pay, upskill and manage hundreds or thousands of knowledge workers, automated mining and reporting of data can be performed by software a fraction of the cost. To allow independent actors, with intentions either good and bad, to exploit this technology unrestricted in a free-market setting, would be to give up the reins to guiding the development of this space in a beneficial direction for all. Sam Altman, president of Y-Combinator and chair of OpenAI, once remarked that

“the best way AI can develop is if it's about individual empowerment and making humans better, and made freely available to everyone.”

BIBLIOGRAPHY

- AI INDEX 2018 (2018) by Stanford University, Yoav Shoham, Raymond Perrault, Erik Brynjolfsson, Jack Clark, James Manyika, Juan Carlos Niebles, Terah Lyons, John Etchemendy, Barbara Grosz and Zoe Bauer, "The AI Index 2018 Annual Report", AI Index Steering Committee, Human-Centered AI Initiative, Stanford University, Stanford, CA, December 2018. AI INDEX 2018 (c) 2018 by Stanford University, "The AI Index 2018 Annual Report" is made available under a Creative Commons Attribution-NoDerivatives 4.0 License (International) <https://creativecommons.org/licenses/by-nd/4.0/legalcode>
- Bulletin of Atomic Scientists. (2018). 2018 Doomsday Clock Statement. Retrieved from <https://thebulletin.org/sites/default/files/2018%20Doomsday%20Clock%20Statement.pdf>
- Choo, Chun Wei (2001). Environmental scanning as information seeking and organizational learning
- Earnest & Young. (2018). What if employment as we know it today disappears tomorrow?: A forward-looking view of the workplace in Germany, Switzerland and Austria in 2030
- Executive Office of the President (2016) Artificial Intelligence, Automation, and the Economy
- Facebook Newsroom (2017), Hard Questions: How We Counter Terrorism, By Monika Bickert, Director of Global Policy Management, and Brian Fishman, Counterterrorism Policy Manager, published June 15, 2017, retrieved from <https://newsroom.fb.com/news/2017/06/how-we-counter-terrorism/>
- Google (2018), Sundar Pichai. CEO Google, Published Jun 7, 2018, Retrieved from <https://www.blog.google/technology/ai/ai-principles/>
- Kurzweil, Ray. (2001). The Law of Accelerating Returns
- Meadows, D. H., & Club of Rome. (1972). The Limits to growth: A report for the Club of Rome's project on the predicament of mankind. New York: Universe Books.
- Rosling, Hans,. Rosling Rönnlund, Anna,. Rosling, Ola. (2018) Factfulness : Ten reasons we're wrong about the world - and why things are better than you think
-

Roser, Max. (2019). Our World In Data. Retrieved from <http://www.ourworldindata.com>

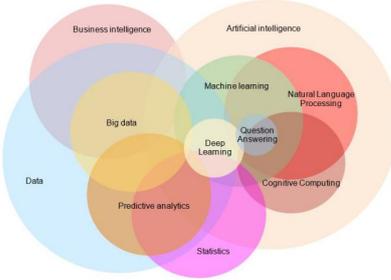
Russell, S. J., Norvig, P., & Canny, J. (2003). Artificial intelligence: A modern approach.

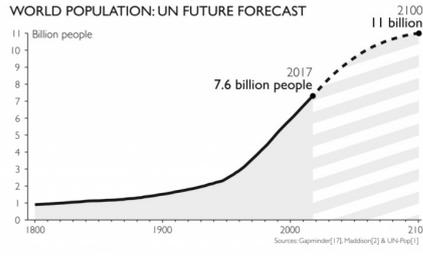
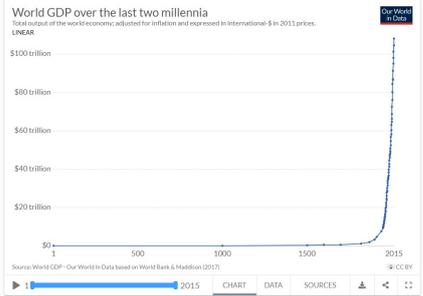
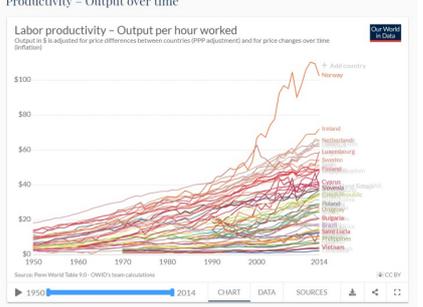
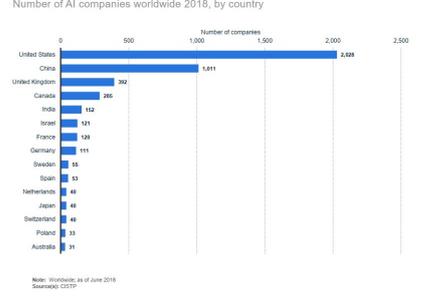
Schwab, Klaus (2016). The fourth industrial revolution., London,. Portfolio Penguin

Statista. (2018). ARTIFICIAL INTELLIGENCE (AI) written by Holst, Arne

Tegmark, Max. (2017). Life 3.0: Being human in the age of artificial intelligence.

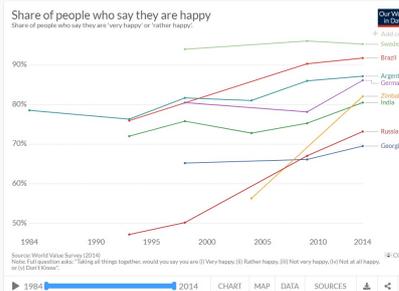
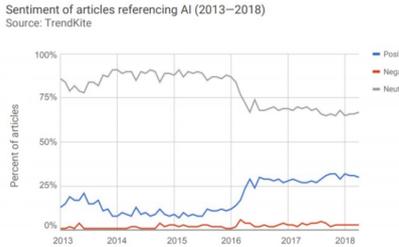
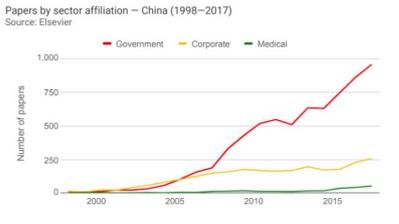
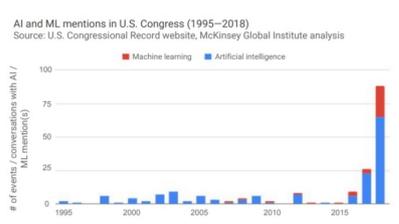
van Vuuren, Rianne. (2016). Information warfare as future South African national security threat

Source	In text citation	Retrieved from:
<p>Figure 1: The many sub-fields and overlapping parallel fields of data science and artificial intelligence.</p> <p>The limits to growth : the 30-year update</p>	 <p>Meadows, Donella H. (2004)</p>	<p>Meadows, Donella H. (2004). The limits to growth : the 30-year update. White River Junction, Vt :Chelsea Green Pub. Co., Retrieved from http://www.peakoilindia.org/wp-content/uploads/2013/10/Limits-to-Growth-updated.pdf</p>
<p>Figure 2. The Doomsday Clock at “Two minutes to midnight” (Bulletin of the Atomic Scientists, 2018)</p>	 <p>Bulletin of Atomic Scientists. (2018).</p>	<p>Bulletin of Atomic Scientists. (2018). 2018 Doomsday Clock Statement. Retrieved from https://thebulletin.org/sites/default/files/2018%20Doomsday%20Clock%20Statement.pdf</p>

<p>Figure 3. The leveling off of the global population as according to Rosling's Gapminder Foundation.</p>	 <p>WORLD POPULATION: UN FUTURE FORECAST</p> <p>11 Billion people</p> <p>2100 11 billion</p> <p>7.6 billion people 2017</p> <p>1800 1900 2000 2100</p> <p>Source: Gapminder[1], The Malthusian[2] & UN Pop[3]</p> <p>Rosling, Hans (2018)</p>	<p>Rosling, Hans,. Rosling Rönnlund, Anna,. Rosling, Ola. (2018) Factfulness : Ten reasons we're wrong about the world - and why things are better than you think https://www.gapminder.org/topics/population-forecasts/</p>
<p>Figure 4. Global GDP over a 2000 year time-frame. (Max Roser, Our World In Data, 2019)</p>	 <p>World GDP over the last two millennia</p> <p>Total output of the world economy adjusted for inflation and expressed in International-\$ in 2015 prices.</p> <p>LINEAR</p> <p>\$100 trillion</p> <p>\$80 trillion</p> <p>\$60 trillion</p> <p>\$40 trillion</p> <p>\$20 trillion</p> <p>\$0</p> <p>1000 500 1000 1500 2015</p> <p>Source: World GDP - Our World In Data based on World Bank & Maddison (2017)</p> <p>Our World In Data</p> <p>2015 CHART DATA SOURCES</p> <p>Roser, Max (2019)</p>	<p>Roser, Max. (2019). Our World In Data. Retrieved from http://www.ourworldindata.com</p>
<p>Figure 5. Labour productivity over time since 1950. (Max Roser, Our World In Data, 2019)</p>	 <p>Productivity - Output over time</p> <p>Labor productivity - Output per hour worked</p> <p>Output per hour adjusted for price differences between countries (PPP adjustment) and for price changes over time (inflation)</p> <p>\$100</p> <p>\$80</p> <p>\$60</p> <p>\$40</p> <p>\$20</p> <p>\$0</p> <p>1950 1960 1970 1980 1990 2000 2014</p> <p>Source: Penn World Table 9.0 - OMD's own calculations</p> <p>Our World In Data</p> <p>2014 CHART DATA SOURCES</p> <p>Max Roser, (2019)</p>	<p>Roser, Max. (2019). Our World In Data. Retrieved from http://www.ourworldindata.com</p>
<p>Figure 6. Number of AI companies worldwide as of June 2018, by country</p>	 <p>Number of AI companies worldwide 2018, by country</p> <p>Number of companies</p> <p>United States 2,029</p> <p>China 1,411</p> <p>United Kingdom 392</p> <p>Canada 285</p> <p>India 112</p> <p>Israel 121</p> <p>France 128</p> <p>Germany 111</p> <p>Sweden 15</p> <p>Spain 51</p> <p>Netherlands 48</p> <p>Japan 44</p> <p>Switzerland 44</p> <p>Poland 33</p> <p>Australia 31</p> <p>Note: Worldwide, as of June 2018</p> <p>Source(s): CRISP</p> <p>(Statista, 2018)</p>	<p>Statista. (2018). ARTIFICIAL INTELLIGENCE (AI) written by Holst, Arne</p>

<p>Figure 7. The share of AI businesses worldwide as of June 2018</p>	<p>Share of AI companies as of June 2018</p> <p>(Statista, 2018)</p>	<p>Statista. (2018). ARTIFICIAL INTELLIGENCE (AI) written by Holst, Arne</p>
<p>Figure 8. Artificial intelligence acquisitions from 2010 to June 2018.</p>	<p>Number of artificial intelligence startups acquired from 2010 to June 2018, by company</p> <p>(Statista, 2018)</p>	<p>Statista. (2018). ARTIFICIAL INTELLIGENCE (AI) written by Holst, Arne</p>
<p>Figure 9. Google's TensorFlow machine learning platform currently leads in the number of developers and projects according to GitHub.</p>	<p>Popularity/usage of artificial intelligence frameworks worldwide 2018</p> <p>(Statista, 2018)</p>	<p>Statista. (2018). ARTIFICIAL INTELLIGENCE (AI) written by Holst, Arne</p>
<p>Figure 10. Share of global AI investment and financing from 2013 to 2018.</p>	<p>Share of global AI investment and financing from 2013 to Q1 '18</p> <p>(Statista, 2018)</p>	<p>Statista. (2018). ARTIFICIAL INTELLIGENCE (AI) written by Holst, Arne</p>
<p>Figure 11. Increase of GDPs globally due to artificial intelligence by the year 2030</p>	<p>AI INDEX (2018)</p>	<p>AI INDEX 2018 (2018) by Stanford University</p>
<p>Figure 12. Chinese company DJI features as the highest valued artificial intelligence company in China in 2018.</p>	<p>Highest valued artificial intelligence companies in China in 2018 (in billion U.S. dollars)</p> <p>AI INDEX (2018)</p>	<p>AI INDEX 2018 (2018) by Stanford University</p>

<p>Figure 13. Active AI start ups in the U.S. increased 113% from 2015 to 2018 while active start ups as a whole increased 28%.</p>	<p>AI startups (U.S., January 1995 – January 2018) Source: Sand Hill Econometrics</p> <p>AI INDEXTM (2018)</p>	<p>AI INDEXTM 2018 (2018) by Stanford University</p>
<p>Figure 14. Venture Capital investment within the US has increased over 16 times its 2010 levels.</p>	<p>Academia-Industry dynamics (2010 – 2017) Source: Sand Hill Econometrics, Scopus, university provided data</p> <p>AI INDEXTM (2018)</p>	<p>AI INDEXTM 2018 (2018) by Stanford University</p>
<p>Figure 15. Michael Huberman and Chris Minns published estimates of weekly work hours going back to the late 19th century. This data shows that over this time working hours have steeply declined.</p>	<p>Weekly Work Hours Work hours of full-time production workers (male and female) in non-agricultural activities. Source: Huberman & Minns (2007)</p> <p>Max Roser, (2019)</p>	<p>Roser, Max. (2019). Our World In Data. Retrieved from http://www.ourworldindata.com</p>
<p>Figure 16. On average, men make up 71% of the applicant pool for AI jobs in the U.S.</p>	<p>Job applicants by gender (2017) Source: Gartner TalentNeuron</p> <p>AI INDEXTM (2018)</p>	<p>AI INDEXTM 2018 (2018) by Stanford University</p>
<p>Figure 17. Gender breakdown of AI professors</p>	<p>Gender break down of AI professors – Select schools (September, 2018) Source: University faculty rosters</p> <p>AI INDEXTM (2018)</p>	<p>AI INDEXTM 2018 (2018) by Stanford University</p>

<p>Figure 18. The share of people either Very happy or Rather happy.</p>	 <p>Share of people who say they are happy Share of people who rate they are 'very happy' or 'rather happy.'</p> <p>Source: World Value Survey (2018) Note: All positive rates. "Taking all things together, would you say you are (1) Very happy, (2) Rather happy, (3) Not very happy, (4) Not at all happy, or (5) Don't know?"</p> <p>Max Roser, (2019)</p>	<p>Roser, Max. (2019). Our World In Data. Retrieved from http://www.ourworldindata.com</p>
<p>Figure 19. Articles on AI became 1.5x more positive from 2016-2018</p>	 <p>Sentiment of articles referencing AI (2013–2018) Source: TrendKite</p> <p>AI INDEX (2018)</p>	<p>AI INDEX 2018 (2018) by Stanford University</p>
<p>Figure 20. Between 2007 and 2017, government AI research papers in China increased by 400% while corporate AI papers in China increased by 73%.</p>	 <p>Papers by sector affiliation – China (1998–2017) Source: Elsevier</p> <p>Note: For reference, in 2017 the number of AI papers on Scopus affiliated with the academic sector in China (not shown) was ~14,000 or ~92% of total AI papers in China.</p> <p>AI INDEX (2018)</p>	<p>AI INDEX 2018 (2018) by Stanford University</p>
<p>Figure 21. AI and ML mentions in U.S. Congress (1995–2018)</p>	 <p>AI and ML mentions in U.S. Congress (1995–2018) Source: U.S. Congressional Record website, McKinsey Global Institute analysis</p> <p>Note: The U.S. Senate and House of Representatives are included in this analysis. Data for 2018 is through 11/20/18.</p> <p>AI INDEX (2018)</p>	<p>AI INDEX 2018 (2018) by Stanford University</p>

<p>Figure 22. Number of components per integrated function for minimum cost per component extrapolated vs. Time. (Log scale)</p>	<p>(Max Roser, 2019)</p>	<p>Roser, Max. (2019). Our World In Data. Retrieved from http://www.ourworldindata.com</p>
<p>Figure 23. Number of transistors which fit into a microprocessor from 1971 to 2017. (Max Roser, Our World In Data, 2019)</p>	<p>Max Roser, (2019)</p>	<p>Roser, Max. (2019). Our World In Data. Retrieved from http://www.ourworldindata.com</p>