



SVNJ Working Paper
2018-2

Artificial Intelligence: The Algorithmic Revolution Driving the Next Industrial Transformation

Kenji E. Kushida, Stanford University

Artificial Intelligence: The Algorithmic Revolution Driving the Next Industrial Transformation

Kenji Kushida

kkushida@stanford.edu

Research Scholar, Asia-Pacific Research Center
Stanford University

Project Leader,
Stanford Silicon Valley – New Japan Project
www.stanford-svnj.org

Prepared for the Trilateral Commission 2018 Plenary in Singapore, March 2018

AI: The New Frontier – and Accelerating

Artificial Intelligence (AI) is rapidly becoming one of the underlying drivers of the next wave of industrial transformations. There is every reason to believe that we are on the cusp of a sea change in how human activities and decision-making are transformed by abundant computing power. This research note will provide the basis for understanding the conceptual building blocks and paradigmatic examples of how the development of AI is accelerating, and how its deployment will be transformative.

What is AI? In Layman’s Terms

The field of AI can be quite broad, and it is not well defined along the edges. Some computer scientists joke that the history of AI is littered with skeletons because areas of research that proved fruitful received their own name (such as machine learning or deep learning), while areas that haven’t had much progress have been lumped into the label “AI”.) The areas that have experienced radical progress recently, leading to a plethora of new possibilities, are generally centered around “deep learning” using “neural networks,” which will be explained below.

It is also noteworthy that some of the confusion around the term “AI” is deliberate. There is commercial pressure to label any analytical service as “AI” since it is likely to sell more than a service that does not -- regardless of what the algorithms are actually doing. Yet, any definition that lacks excludability ceases to have any significance.¹

¹ Silicon Valley accelerator Y Combinator even has a blog entry titled, “how to know when products actually use AI.”

<http://blog.ycombinator.com/how-to-know-when-products-actually-use-ai/>

As a general principle, for any phenomenon for which we use a label, we need a definition that includes the ability to exclude things; whenever a word becomes a buzzword, such as “Cloud-based,” “Big Data” or “AI”, the term itself loses meaning unless one can say which things are “not AI”. Since all IT services and data analytics are now under marketing pressure to label themselves “AI” or “Big

For our purposes, **a simple working conception of AI centers around pattern recognition—such as finding underlying relationships in data, images, audio, video—that is applied to enable a software system to improve its capabilities, or “learning.”**

Self-driving cars, for example, are learning about their environment through recognizing patterns of visual data, sometimes combining them with other data, such as maps. *Natural language processing*, for example, relies on recognizing and learning about language patterns to correlate with meanings, and audio input relies on sound pattern recognition of spoken words. *Motion and manipulation*, which is critical for robotics, often employs a variety of visual and tactile data inputs to enable the system to learn about its environment and how it interacts with it. A nice example of the latter is a little humanoid robot learning to swing on a mini swing-set. A conventionally programmed robot would need extensive programming about when to move its legs at what timing in order to amplify the center of gravity changes that enable it to swing, and then how to adjust as the swinging got bigger. It took a couple days to program such a robot, with numerous adjustments along the way. The same robot that used machine learning had no information about how to operate a swing set and only knew how to swing its legs, and the objective of attaining a pendulum motion. It first randomly swung its legs until it learned which timing would best amplify the swinging, and rapidly improved itself. In a matter of hours, it seemed to “get the hang” of how to swing, and began swinging vigorously – with no human interaction.

The easiest, though simplest way to determine if something deserves to be called AI is to inquire **what is the pattern being recognized from what data**, and **how is the learning occurring and feeding back** to the pattern recognition ability.

Some Vocabulary for Basic Concepts

The field of AI covers a variety of approaches to solving problems and learning, and it is worth becoming familiar with some of the basic terms. Many of these approaches and tools were developed decades ago, but have recently undergone breakthrough improvements due to the availability of vastly increased computing power.

Machine learning is a broad categorical term, but specific types of approaches include probabilistic tools, for when the information available is incomplete or uncertain, classifiers and statistical learning models, such as placing certain observations in particular categories, and logic approaches for when learning is sequential.

Deep learning is a methodology employing *neural networks*, patterned after how the human brain works, and is used widely in current pattern recognition. There are multiple layers of relatively simple calculation between the input, which is known, and output, which can be observed and assigned a score.

There are also several types of learning.

Supervised learning teaches patterns for which the right answer is known. For an example of supervised learning, Facebook or Google have millions of images of cats, for which people have labeled cats in their photo albums. The AI program is then given an

data analysis,” we have to be sure what this excludes. Suppose we have an extremely effective analytical algorithm that actually does not use a great amount of data, and it is a conventional algorithm – it would be highly unattractive if labeled, “small data, non-AI analytics.” Thus, everything is under pressure to be called “AI” and “big data.” However, this obscures the true nature and therefore the extraordinary potential of AI.

image, and must try to determine whether or not it is a cat. The program's success or failure is known, and it can improve. We do not know, however, the basis of its decision to identify a cat—shape of nose, ears, etc. A traditional algorithm would have to start by specifying parameters for a cat, such as pointy ears, then apply each picture to the rules. The exceptions then become problematic – what if the cat has partially chewed up ears, or is wearing a hat, for example, so that has to be taken into account ahead of time. With AI, however, the algorithm learns by itself what to look for without being told.

Unsupervised learning entails simply teaching the AI engine the “rules of the game” and letting it find underlying patterns or categorizations. For example, an unsupervised learning program to play “Go” will know only the rules, and whether it wins the game or not, without using historical data from past games. Its first games will therefore be random, and it will lose badly. After many iterations, however, it will figure out the underlying patterns that are likely to result in losing less badly, then eventually winning more and more decisively.

The term Artificial “Intelligence” is sometimes misleading, since the word “intelligence” can conjure images of computers developing a “consciousness,” or achieving “*Singularity*,” a commonly used word with somewhat unclear meanings, but generally in which computers exceed human capability in all capacities. There are several concepts that must be sorted out to have a reasonable discussion. On the one hand, almost any technology is used by humanity precisely because it outperforms humans at certain tasks. A top fighter jet pilot discovered that once she became an elite of the elite, the planes she flew were so automated that she was not allowed touch the controls—the error margins for takeoff and landing on an aircraft carrier were so narrow that only computers could perform these tasks.

On the other hand, it is the specter of a single general purpose program or computer that can outperform humans on *all* different tasks that humans perform that often scares people. In the fighter jet example above, this would be the onboard computer also cooking the pilot's breakfast, filing her paperwork, making personnel decisions about those under her command, and managing her parents' bank transactions. We are still far away from the latter.

Yet, the notion of *consciousness* is very different from attaining skills to perform tasks, which should be treated as a different dimension from performance. Researchers at the forefront of AI generally consider matters of consciousness or singularity to be far off enough in the future and unclear enough in how such phenomena will unfold, that we are a long way from worrying about these developments.² (That being said, many industrial disasters and most disaster movies start out with scientists and engineers vastly underestimating the power and unintended consequences of their creations.)

² It might be noted that a prominent pioneer of autonomous vehicle driving, Anthony Levandowski, head of Google's autonomous driving project before leaving to set up his own company, which was then acquired by Uber, subsequently left Uber after a lawsuit filed by Google, and has set up his own church to worship the “godhead of AI” – but this probably deserves a footnote treatment rather than a serious discussion for now. <https://www.wired.com/story/anthony-levandowski-artificial-intelligence-religion/>

Why is AI developing so Rapidly Now?

The speed of recent developments in AI is surprising not only to the general public, but to many specialists in the field. The development is driven largely by an underlying civilizational transformation in the basic resources available to our civilization—computing resources. ***Computing resources—the ability to compute, store, and transmit information—has recently transformed from a scarce to an abundant resource for the first time in human history.***³

Throughout human history, computing resources have been scarce, and therefore costly. From stone pyramids through the invention of the atomic bombs, most of humanity's complex calculations and mathematics were done largely by hand.

The extreme nature of exponential growth in processing power following the invention of the semiconductor is often underappreciated. As Intel founder Gordon Moore's development objective of doubling the number of transistors on a semiconductor chip every 18 months (Moore's Law) held from the mid-1960s onward, computational power available to humanity has grown at astonishing rates.

For example, the processing power available to the main computer in the Apollo mission to the moon in 1969—a major milestone in human civilizational attainment—was roughly equivalent to the processor in the Nintendo Family Entertainment system that debuted in 1983, retailing for under \$100 and marketed to children.

The fastest supercomputer in 1985, the Cray II supercomputer, had a processing power roughly equivalent to one sixth that of the iPhone 6, introduced in 2014. Yet, while there were only a handful of Cray II supercomputers in existence in 1985, the number of smartphones shipped in 2017 was over 1.5 billion.⁴ *Most people in the developed and developing world are carrying around far greater processing power than the world's fastest supercomputer 30 years ago.*

A 2016 Intel processor compared to a 1971 Intel processor had 3500 times the processing power, 90,000 times the energy efficient, and 1/60,000 the price. Intel engineers (with a good sense of humor) calculated the equivalent performance increase for a 1971 Volkswagen Beetle. *If the performance of a 1971 Volkswagen Beetle improved at the same rate along the same dimensions as semiconductors, the 2016 model would have a top speed of 2800 miles per hour, the fuel efficiency would be such that one gallon of fuel would allow it to travel 2 million miles, and the price would be 4 cents.*⁵ Of course, the late 2017 model, debuting 18 months later, would double this, traveling at 5600 miles an hour, 4 million miles per gallon, and 2 cents.

This is the driver of computing power abundance.

AI programs and experiments that used to take weeks and days to complete can now be performed in milliseconds. Many of the underlying theories and concepts in AI date from the 1950s and 60s, but it was only recently that enough processing power could be mobilized at low enough costs to solve problems and discover new methods.

³ Kushida, K. E., J. Murray and J. Zysman (2015). "Cloud Computing: From Scarcity to Abundance." Journal of Industry, Competition and Trade **15**(1): 5-19.

⁴ <https://www.statista.com/statistics/263437/global-smartphone-sales-to-end-users-since-2007/>

⁵ Friedman, T. L. (2016). Thank you for being late: An optimist's guide to thriving in the age of accelerations, Farrar, Straus and Giroux.

Complementary Enabling Technologies: Cloud Computing, Sensors, Smartphones

Throughout the history of technological change, it is rarely the development of a single breakthrough technology that shapes its diffusion throughout the world. Rather, *it is when complementary technologies are implemented, each through their own market or industry dynamics, that the technologies in question reveal their potential.*⁶

Historically, for example, the invention of asphalt, critical to driving automotive vehicles smoothly, was a necessary complementary technological innovation to unlock the potential of the internal combustion engine as implemented in cars. Once it was technically possible to build smooth roads cheaply, it was then it was up to the political decisions to build paved road networks that enabled trucking to displace trains as a low-cost method for moving freight.

For AI, the set of complementary technologies include global scale cloud computing, the advent of low cost sensors, and the vast diffusion of smartphones.

The radical gains in computing power enabled by Moore's law in transistor development became available broadly beyond the exclusive domain of leading edge firms with massive datacenters, due to the advent of **global-scale Cloud Computing**. Only then could every startup and researcher with even relatively modest financial resources harness the power of civilization's most powerful computing infrastructure. A handful of companies including Amazon, Google, Apple, and Microsoft, are able to offer low cost computing through massive datacenters around the world, constantly updated with the fastest physical resources and software enabling ever-more efficient use of the raw processing power. The Cloud offerings enable users, firms and individuals, by providing low cost access to the frontier of abundant processing power.⁷

Another critical ingredient for the leaps and bounds in AI development has been the availability of massive amounts of data to feed the algorithms. The **diffusion of low cost sensors** to measure a variety of things has been indispensable to creating the vast data that is most useful for AI.

The diffusion of **smartphones** has been a driver of diffusing sensors throughout the world. An ordinary smartphone contains numerous sensors, such as accelerometer, proximity sensor, light sensor, barometer, thermometer, gyroscope, magnetometer, pedometer, air humidity sensor, and of course camera, GPS for location, and microphone for sound. The price of these sensors had dropped dramatically – almost 100 fold for some that are included in smartphones due to advances in nanotechnology and the sheer volume of smartphones shipped.

The sheer quantity of data that can be collected cheaply now enables companies such as Alphabet to use the GPS of Android smartphone users to deliver the real time traffic information on Google Maps, for example.

In various areas of industry, the **Internet of Things (IoT)** is about how to collect good data and measuring things that have not previously been measurable or observable. With the data fed into AI models for machine learning, pattern recognition and predication can become valuable tools for managing various aspects of industry.

An important perspective moving forward is how to effectively collect the data you want in clever ways by using sensors, and then using AI tools to identify

⁶ Perez, C. (2009). "Technological revolutions and techno-economic paradigms." Cambridge journal of economics: bep051.

⁷ Kushida, K. E., J. Murray and J. Zysman (2015). "Cloud Computing: From Scarcity to Abundance." Journal of Industry, Competition and Trade **15**(1): 5-19.

patterns, feed those outcomes into services, products, or other value-added offerings.

For example, if one airline experience is consistently far better than another's, but they are engaged in price competition, can this difference in service quality be measured objectively through biometric data from passengers willing to wear simple devices in exchange for airline miles? The cost for such an experiment, with AI mobilized to find underlying patterns, is no longer prohibitive.

The Impact of AI: How will it unfold?

We are only at the beginning of seeing the real impact of AI.

As with many moments of remarkable technological leaps that become the underlying building blocks for vast advances in industry, we are likely to see the impact of AI unfold in several different ways. Here it is useful to distinguish between **frontier companies** that are at the forefront of pushing the frontier of AI forward; they bring massive quantities of data unmatched by others, vast financial resources, and the ability to attract top talent to pioneer projects that go beyond considerations of immediate commercialization. Most companies are not at the frontier, and their relationship to AI is to use the tools, services, and platforms offered by frontier companies.

Specialized tool and industry AI companies use smaller data than frontier companies and often utilize the computing resources or tools provided by frontier companies, but lead the way in applying AI in specialized areas of problem solving, or for particular industries. Examples are below.

Ubiquitous AI will be a phase when AI becomes so pervasive that it provides the baseline to the extent that it becomes nearly invisible—almost as computers are today. Just as companies now do not boast about using computers in their daily operations, or have Internet connectivity, it will not only make little sense to advertise that they use AI, but it will even be difficult to distinguish which tasks do not rely on AI at some level in the various layers of computing infrastructure and services.

Frontier: A Handful of Companies with Vast Data and Financial Resources

Frontier companies will be pushing the forefront of knowledge on various forms of deep learning, using massive amounts of data and processing power that are unavailable elsewhere. Only these companies, with the most data and financial resources, will be able to hire truly top talent, who will continue to push the frontier forward. Currently, only a small handful of companies are in this position – Alphabet (owner of Google), Amazon, Facebook, and Microsoft.

GAFA – Google, Amazon, Facebook, Apple, are currently the destinations of choice for many graduates of top notch computer science programs, such as Stanford. The amounts of data available to these companies are an order of magnitude greater than those held elsewhere. An interesting study notes that in hiring, places such as IBM are unable to hire specialists in deep learning because the industrial data available to companies such as IBM are far less than those available elsewhere.⁸

The frontier companies also have not only among the highest market capitalizations in the world, but also the largest cash piles.⁹ They are able to hire top

⁸ <https://techcrunch.com/2017/07/13/jefferies-gives-ibm-watson-a-wall-street-reality-check/>

⁹ <http://www.businessinsider.com/chart-us-companies-with-largest-cash-reserves-2017-8>

talent, and undertake ambitious projects. Such projects do not need to have immediate commercial application, but these firms can employ researchers to push the frontier forward.

The most paradigmatic example of frontier company activity is Deep Mind, the company purchased by Google. Deep Mind created AlphaGo, a program to play the board game of Go, long considered too complex for algorithms to beat humans.¹⁰ In 2015, AlphaGo beat a human professional for the first time, and in 2016, AlphaGo famously beat a top professional, Lee Sedol, in 4 out of 5 matches.

However, so what if AI can play go; what does it mean for the real world? In a news release that surprisingly gained far less attention than it deserved, in the summer of 2016, Google used Deep Mind to optimize the cooling of its datacenters. According to Google's estimates provided to a Stanford researcher, the company consumed 0.01% of the world's electricity in its massive datacenters in 2011.¹¹ Optimizing cooling is therefore an important task to reduce electricity consumption. Using a program written by Deep Mind, Google found that it was able to increase the efficiency of its datacenter cooling by 40%, reducing electricity consumption by 15%.¹² This was a significant application of AI with tangible real-world benefits.

The critical question for our current society is when Google will make these types of AI tools for practical application available to the general public. For now, Google's method is only available within the company, and we do not know how much processing power was consumed to optimize its datacenter cooling, and how much data was needed, or what types of data were used. However, ***it is when the cutting edge technology offerings become commonplace commodities that vast new possibilities open up.*** It is quite possible to imagine programs with industrial applicability by Deep Mind, or similar AI offerings, to become externally facing for extremely low costs—perhaps a subscription of perhaps \$10 or \$20 a month.

Amazon, Microsoft, and Google have already begun opening some of their AI engines up to the public, some with tools that enable non-experts in machine learning to use machine learning tools. When an array of powerful tools like this becomes available to companies more generally, the full impact of AI will begin to be felt.

The vast majority of companies are not frontier companies, but rather users of the tools that frontier companies provide. The question then becomes ***how prepared and how effectively follower companies can use these tools when they become available.***

As tools become available, they will become the *new baseline for cutting costs*. This type of implementation will be necessary to compete, but there is likely to emerge

¹⁰ In 1997, an IBM computer, Deep Blue, had been the world champion chess player Gary Kasparov, but Go is a far more complex game, and computer programs were unable to beat top ranked professionals without handicaps until AlphaGo.

The Economist provides an excellent description of the difficulty of Go. "A 19x19 board offers 361 different places on which Black can put the initial stone. White then has 360 options in response, and so on. The total number of legal board arrangements is in the order of 10^{170} , a number so large it defies any physical analogy (there are reckoned to be about 10^{80} atoms in the observable universe, for instance).

<https://www.economist.com/news/science-and-technology/21730391-learning-play-go-only-start-latest-ai-can-work-things-out-without>

¹¹ <http://www.datacenterknowledge.com/archives/2011/08/01/report-google-uses-about-900000-servers>

¹² <https://deepmind.com/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-40/>

a robust industry that offers such tools to corporations, making cost cutting alone unlikely to be a competitive differentiator. *The question is what value-added activity can be created from the tools.*

At the same time that frontier firms are pushing the boundaries of AI, we are seeing an explosion of companies offering specialized tools and industry applications.

Specialized Tools and Industry Applications

The vast data, talent, and financial resources of the leaders enable another set of firms to offer **specialized AI tools and industry applications** as they are able to use resources offered by the frontier firms and make use of the advancements that frontier firms enable.

Cloud computing services such as Amazon's AWS, which allows users to pay for the amount of computing power they use with virtually no capacity constraint, and Google's offerings such as databases that can handle extremely large amounts of data, all at low cost, have allowed firms to utilize the world's most powerful datacenters at extremely low costs. This enables companies developing specialized AI tools to flexibly scale up their use of computing power, easily run experiments, and offer their own services built on top of these tools.

As the frontier companies develop more and more machine learning and deep learning tools, along with other sources of data available, the ability of specialized AI tools and industry application companies to make their own offerings is accelerated.

Let us run through some examples – some real world and some within the technological realm of possibility that now simply requires some enterprising companies to test and implement them.

Insurance: Dynamic Pricing and Measuring Behavior Directly

Automobile insurance is legally required for people owning cars in most countries. Yet insurance premium rates are most often calculated through proxies. For example, the type of car you drive, your age, amount of driving, whether you use the vehicle to commute to work, your income level, education level, prior history of traffic violations, etc. However, we are now at the cusp of being able to measure peoples' driving directly—through sensors attached to cars, with information sent by smartphones to networks.

The finer the granularity in capturing data about how people drive, the more sophisticated the models that can be created by feeding driving pattern data into AI algorithms to determine the level of riskiness of drivers.

Rather than using current proxies, direct measurement of human behavior, collected at massive scale, can be analyzed using AI.

For auto premiums, one could also add dynamic pricing—those with risky driving despite appearing to be good drivers in the proxies, could see their premiums rise. By driving well, premiums could decrease. Then, if this incentivized people to drive more safely, there would be fewer accidents, which should benefit insurers as well as society more generally.

The barrier to adopting such models by incumbent insurers is often the legacy IT systems that were not built to cope with the massive amounts of data input and processing that can be done with frontier firms' resources. This opens room for startups or hungry second or third tier firms to aggressively attempt to disrupt the leaders.

Medical Image Diagnosing

In November 2016, Google published a paper in the Journal of the American Medical Association showing how AI could be used to recognize diabetic retinopathy, an eye disease that can lead to blindness if not detected early. Google used image recognition, similar to identifying and labeling people and objects in pictures. Deep Mind was partnering with the National Health Services of Great Britain to identify a variety of diseases and ailments.¹³

In November 2017, Stanford University computer scientists teamed up with Stanford hospital doctors and published a paper documenting the results of a new deep learning algorithm: the AI algorithm outperformed doctors in detecting pneumonia from patient X-Rays. Pneumonia is notoriously difficult to diagnose. The rapid progression of deep learning results is documented in the Stanford Report news: “Within a week the researchers had an algorithm that diagnosed 10 of the pathologies labeled in the X-rays more accurately than previous state-of-the-art results. In just over a month, their algorithm could beat these standards in all 14 identification tasks. In that short time span, CheXNet also outperformed the four Stanford radiologists in diagnosing pneumonia accurately.” (Stanford Report Nov 15, 2017.)¹⁴

“Fintech” (Beyond the hype): New ways of Measuring Risk, Customization

Measuring risk is a core function in finance. Finance was one of the early adopters of computers to perform calculations—especially in insurance.¹⁵ More recently, the financial sectors have driven productivity gains, particularly in the US, through massive investments in IT.¹⁶ The financial crisis of 2007-2008 revealed the extent of risk calculations that were conducted by software, as the vast array of complicated risk assessment models discovered that they had actually severely miscalculated how correlated their risks were, based on assumptions such as continuously rising real estate prices, which did not hold.

There is a broad range of categories of services considered *Fintech*, including personal lending, microfinancing, equity financing, personal asset management, virtual currencies and exchanges, credit rating, and others. A key theme for many of these areas is *new ways of measuring risk, which use data previously unavailable or uncorrelated with individuals, using machine learning to improve credit risk calculations.*

Using data to analyze peoples’ relative risk is nothing new. Almost 30 years ago, word on the street was that American Express used purchases to forecast potential credit risk—specifically, if a married man suddenly bought flowers and jewelry, this suggested an extramarital affair, raising the risk of a divorce, which would often damage him financially.

Today people’s behavior can be analyzed in new ways and correlated. The Chinese firm Tencent, which operates WeChat, one of the two Chinese leading platforms for an array of IT services, including payments, is a leader in using human behavior data across domains to assess risk. As explained by a Tencent executive, for a WeChat user to

¹³ <https://www.wired.com/2016/11/googles-ai-reads-retinas-prevent-blindness-diabetics/>

¹⁴ <https://news.stanford.edu/2017/11/15/algorithm-outperforms-radiologists-diagnosing-pneumonia/>

¹⁵ Cohen, S., J. B. DeLong and J. Zysman (2000). Tools for Thought: What is New and Important about the "E-conomy". Berkeley, CA, Berkeley Roundtable on the International Economy, University of California at Berkeley.

¹⁶ Jorgenson, D. W., M. S. Ho and K. J. Stiroh (2005). Information technology and the American growth resurgence. Cambridge, Mass. ; London, MIT.

borrow one of the ubiquitous rental bicycles, the system assesses the risk score of each individual. If someone has a propensity to gamble, or likes to drive cars late at night at high speeds, their risky behavior lowers their credit score. If the credit score is too low, they cannot rent the bicycle.

Automated Driving: The Holy Grail

There is currently an investment rush into Silicon Valley by existing automobile manufacturers to develop autonomous vehicle driving capabilities. Pattern recognition is the critical task. Google had raised major awareness of the technological possibilities after it hired Sebastian Thrun, who had directed the Stanford Artificial Intelligence Laboratory and won the DARPA's 2005 Grand Challenge with a robotic vehicle. Google's first autonomously driven vehicle was licensed in Nevada in 2012, and it made headlines when it debuted a prototype without a steering wheel, accelerator, or brake pedals—only an emergency stop button. This experimental vehicle fleet could be seen driving around Mountain View near Google headquarters regularly.

Tesla delivered a paradigm shift when it offered its Autopilot in October 2014. An email from Elon Musk noted that existing customers of its Model S sedan could download the program to their cars overnight, transforming the human activity of driving overnight. The initial version of Autopilot was so effective that it raised concern among regulators when videos surfaced on Youtube made by drivers showing themselves driving with no hands on the steering wheel, even playing an instrument, eating a meal, or taking a nap (pretend or real). A software upgrade required drivers to keep their hands on the steering wheel for the majority of the time that Autopilot would be engaged.

Uber sent shockwaves to the community when it announced suddenly in December 2016 that it would begin deploying a limited number of Uber self-driving cars in San Francisco.

Google noted that it spent over \$1 billion USD on its automated vehicle project between 2009 and 2015, with the division spinning out into its own company, WayMo, in 2016. In 2016, GM purchased Cruise Automation for \$1 billion, and Ford created a joint venture with Argo AI, which was only two months old at the time and had been founded by a former Google autonomous vehicle engineer. In 2016, Uber purchased the startup Otto, a six month old startup developing systems to enable semi-trucks to drive autonomously, created by Anthony Levandowski, who had left Google's autonomous vehicle project, for almost \$700 million in 2016. In 2017, Intel purchased an Israeli startup that provided the original vision-based technology to Tesla for its autopilot, for \$15.3 billion USD.¹⁷

When fully automated driving becomes a reality, its effect on lifestyles and patterns of mobility and settlement will be profound – but it is still unclear in what ways. For example, some people could choose to live far away from their work, using their time to sleep or work, but fully automated driving could also enable vehicles to drive much faster than now. This would actually increase their energy consumption in transportation and potentially accelerate urban sprawl. On the other hand, given that almost of a fifth of American cities devote their space to parking lots, automated driving—especially if car sharing allows private vehicles to operate during work hours to carry short distance transportation customers—could play a role in dramatically

¹⁷ <https://spectrum.ieee.org/cars-that-think/transportation/self-driving/google-has-spent-over-11-billion-on-selfdriving-tech>

increasing the comfortable density of urban areas, enabling more people to live close to their work. This would increase energy efficiency and urban density.

Material Informatics: Examples of Utilizing Big Data and Analytics

The number of specialized materials in our world has multiplied radically through advances in materials science and the ceaseless labor of large firms to create new materials for specialized uses. Think of the thousands of specialized materials used in modern cars, compared to the six or seven basic materials that were used in the 1960s.

The creation of new materials for specialized uses is a labor-intensive process within large companies, in which particular properties of the resulting material are desired, but companies need to form hypotheses, then test thousands of times to get results they want.

In the area of alloys, for example, there are about 28 significant parameters, and to find the desired output, scientists must effectively grapple with optimizing the values across these 28 parameters. Of the thousands of experiments, only those yielding the desired properties are considered a success, and the others are often put aside and stored in file cabinets, spreadsheets, and pdf files in a decentralized manner without later utilization.

Citrine Informatics, a Silicon Valley company, pulls all this information from a variety of formats into a database, and provides information for companies to conduct new experiments.

The service fills in the parameters for values for data obtained from the company, and for parameters lacking data, it uses machine learning algorithms to create estimates. Therefore, when a company specifies the properties it wants, Citrine is able to provide a recommendation based on real data the company has generated previously for some of the parameters, combined with estimates about the other parameters based on its own AI machine learning model.

Then, when large company conducts an experiment based on the recommendation by the software, the AI-estimated values can be replaced by real values. The real values are then fed into the machine learning algorithms to improve its quality for subsequent experiments.

Citrine was founded by a computer science PhD in machine learning, and a Stanford business school graduate with an advanced degree in materials science.

This pattern of using large datasets gathered from large companies, filling in parameters with values from the data, and using machine learning to estimate other values, with experiments allowing for improvements to the machine learning algorithm, is a promising new area of AI applications in industrial uses.

Startups using the same principles in agriculture for chemical or “organic” fertilizers also reveal the promise to various areas of economic activity.

AI in Retail: Logistics Optimization and Understanding/Shaping Human Behavior

Retail is a significant area for both the application and development of AI. *Optimizing complex logistics, predicting human behavior based on a variety of behavior and environment data, and discovering underlying patterns in human activity* are all strengths of AI, given the collection of vast and valuable data.

IT has long been used in retail to optimize logistics and make predictions—for example, the convenience store chain 7-11 in Japan has used weather reports to build

models and automatically adjust order quantities and selection of items such as perishable lunches (as well as umbrellas, of course.)

German retailer Otto famously used an AI deep learning algorithm, originally developed for particle physics experiments at CERN, to predict customer purchases a week in advance. The reported 90% accuracy of the system's predictions for purchases within 30 days emboldened Otto to commit to purchasing 200,000 items per month automatically from third party brands – already as of a year ago, in April 2017.¹⁸

Moreover, massive, complicated logistics operations are not only opportunities for companies to deploy AI in order to improve efficiency, but they are also opportunities for frontier firms to gather large amounts of data that can then be mobilized to develop further AI tools—not just for retail, but in other areas as well.

Amazon's purchase of Whole Foods, a high end supermarket, should be seen in this light. Amazon has been a leader in capturing human activity, both in its own warehouses, but also through its wide online retailing operations. Rather than simply aiming to improve the efficiency of the retailer Whole Foods, Amazon's goals clearly included access to incredible amounts of data from a vast and complex supply chain that it can apply in its operations elsewhere. Moreover, information about consumer behavior, if correlated among multiple data sources, can yield insights into potential demand as well.

Much of current retail optimization in IT uses information about what is sold to predict consumer behavior. The question is how to generate data about what people did not buy, but would have bought had it been available. *Retail models often use data about what was sold to find correlations, but if tools such as low-cost voice recognition can be mobilized in smaller stores to identify what local customers would have bought but could not find, even more fine-grained, locally optimized models based on actual demand could be possible.* In this case, AI enters as a tool for voice recognition (since it is unrealistic to expect cashiers to enter information about what customers wanted but could not find manually), and a powerful tool for finding patterns.

AI and Security

Security is a broad area in which AI is beginning to play a role, which will likely gain rapid significance. There are several types of security and it is worth differentiating between *cybersecurity, which is a concern about data—preventing unauthorized access or tampering with data*, and making sure the IT systems are operational at all times to allow both the availability of data and *data privacy, which is about the rules and regulations surrounding the data of individuals.*

The use of *AI in physical security* in applications such as recognizing people as they appear in various places as video images captured in cameras of all sorts, and their digital footprints as they leave traces in financial and online activities, are part of identifying and tracing people. For example, the American company Shot Spotter uses a variety of sensors and AI in urban areas to triangulate gunshots when they occur, since most gunshots are not reported to the police. A listed company, Shot Spotter operates in over 90 cities around the world, including Cape Town, South Africa among their non-American customer cities.

Companies such as Hikvision, from China, operate video surveillance equipment and analyze data, using it to learn about risks. Unattended bags in crowded venues, for

¹⁸ <https://www.economist.com/news/business/21720675-firm-using-algorithm-designed-cern-laboratory-how-germanys-otto-uses>

example, and the behavior of people preceding crimes caught on camera provide input to teach the algorithms to look for potentially risky behavior.

Silicon Valley company Palantir has built itself as a leader in identifying patterns of transactions and various data, ranging from financial data to the movement of people, to conduct surveillance and identify national threats—although the full scope of their activities is classified.

Using *AI to identify people who are potential criminals or terrorists before they commit a crime* is an area where privacy concerns are weight with security considerations, leaving many of the various implementations to national policy arenas. Data available in some countries or regions is unavailable in others to use for law enforcement. For example, Tencent uses a variety of online behavior surveillance tools to identify people who are visiting illegal gambling sites, which are often masked as advertisements or other services, and which frequently move around when detected. Chinese company Cloud Walk Technology is an example of a company that actively tries to identify potential criminals through behavior analysis in security cameras and correlated purchases. The US Federal Bureau of Investigation has long used credit card records and other data inputs to identify potentially risky behavior by identifying people who purchase goods that have a high correlation to terrorist activity, such as bomb-making materials, but the hope is that AI will drastically improve their effectiveness.

An emerging area called **RegTech**, “Regulation Technology,” includes areas such as fraud detection and companies’ compliance with regulations. For example, for large financial institutions, detecting financial fraud and quickly reporting it to government is an extremely costly activity, which entails significant fines if not complied with. A range of startups have appeared which take companies’ financial data, analyze it using machine learning algorithms, and report whether they detect any fraudulent or irregular activity.¹⁹

AI in Agriculture: Agritech

Agriculture is an area ripe for the application of AI, coupled with, and enabled by low cost sensors capable of gathering data and learning about various correlations. While agriculture in areas with large landmasses such as the US has been pursuing economies of scale to gain efficiency, most agricultural sectors in the world are unable to do so. Moreover, variations in the nutritional content in soil, nutritional needs of vast varieties of plants, complex logistics and supply chains, and as a sector with relatively low skilled labor, the field of “Agritech” is receiving significant attention.

Data gathering can be done more effectively than ever before, through means such as drones and collecting data using the agricultural machines themselves. The chemical compositions of soil and their incredible variations, and of organic fertilizers, for which the molecules cannot be altered like those of chemical fertilizers, are areas in which platforms like the material informatics service above, can be useful.

Predicting rainfall, cold snaps, or other extreme weather conditions can be critical to crop yields. Accurate weather and climate prediction models become of paramount importance, especially at the very micro-level. This is a computationally intensive challenge, and the use of low cost sensors to get granular data to build models is critical, and can be improved by AI.

¹⁹ For more, see the CB Insights report on Regtech, although they are more narrowly focused on the financial sector. <https://www.cbinsights.com/research/regtech/>

Mechanization and increasing labor productivity is also a challenge for agriculture—another core domain for AI. Understanding the tasks that need to be performed at short notice, e.g. when high value crops such as fruits ripen, requiring bursts of labor. How to best mechanize difficult tasks like picking raspberries or easily bruised pears which require attaining speed and scale are among the challenges that the current state of AI and robotics, along with logistics optimization, is ready to tackle.²⁰

The AI Frontier Keeps Moving: Less Data Needed?

As has been shown so far, advances in AI have been enabled by increases in processing power availability and capabilities to gather massive amounts of data for deep learning.

Recently, frontier research suggests the potential for far small amounts of data required for effective learning. In a paper published in *Nature* in October 2017, Google announced that it had created a new AI program called AlphaGo Zero. The original AlphaGo had employed “supervised learning,” in which it learned to play Go by analyzing data of a large number of previous Go matches to learn tactics and strategies, then engaged in “reinforcement learning” by playing against itself to improve its skill. AlphaGo Zero, on the other hand, was engaged in “unsupervised learning” and was only taught the rules, and engaged in “reinforcement learning” by playing itself. While AlphaGo played itself approximately 40 million times over the course of several months, AlphaGo Zero played itself only around 1.5 million times.²¹

When they played each other AlphaGo Zero beat the original AlphaGo by 100 matches to zero. This is the speed at which AI is improving.

However, the new power of unsupervised learning is still only in the realm of games. Yet, past experience suggests it would be premature to consign this type of learning only to the realm of games.

Frontier firms continue to absorb top talent. For example, a leading AI researcher, Fei Fei Lee, was recently hired from Stanford University by Google, and the latter announced that she would be in charge of creating an AI lab in China—despite Google not having core operations there.²²

Technology Diffusion Relies on Context

For any technology, the pattern in which it diffuses is determined not primarily by attributes of the technology itself, but by the context. The context includes related complementary technologies, industry dynamics, regulations, politics, and other social factors such as relative cost and availability of labor.

In a historical example, the technological innovation of steam power had a globally transformative impact through the railroad, transoceanic shipping, and factories, which reshaped global trade, production systems, the migration of people, and the accompanying reshaping of global civilization. The complementary technology of advances in steel making (Bessemer steel) rose out of an industrial context, but canals and the advent of shipping lanes, railroads, and global systems of trade all arose from

²⁰ For a list of interesting agritech startups, see the participants in Nikkei’s Agritech Summit (Ag/Sum). <https://www.agsum.jp/en/exhibitor>

²¹ <https://www.nature.com/articles/nature24270>

²² <https://www.blog.google/topics/google-asia/google-ai-china-center/>

political decisions about how to use resources, and how nations chose to interact with each other.

AI, which necessarily involved the use of data, with much of the data being from people who live in countries with varying sets of rules, will be shaped by national and regional contexts.

Contemporary real world examples of “rules matter” in the deployment of AI are obvious. For autonomous driving for example, Nevada passed a law in 2011 for autonomous cars, requiring a person in the driver and passenger seats during tests. Google’s car was the first autonomous vehicle licensed at the Department of Motor Vehicles in 2012. When Uber suddenly announced its deployment of self-driving vehicles for passengers in 2016, The California Department of Motor Vehicles contended that Uber had not obtained proper licenses, and revoked the car registrations. Uber then moved its testing to Arizona, which welcomed them with open arms, in a highly publicized strategy of regulatory arbitrage.

Tencent’s ability to use various personal data to build models of people’s credit risk is allowed by Chinese law, but privacy laws in the EU often forbid personal data from crossing national borders. While there are good arguments for this, the question is whether it creates a disadvantage in developing and deploying systems that can benefit from AI.

For any technology involving human genome alternation, rules and regulations will certainly matter.

Rules and regulations, of course, do not exist in a vacuum, but are the product of political processes, which differ across countries. Powerful industry groups, voters, the media, other political agendas, and the bureaucracies are some of the key actors that are empowered very differently across different political regimes around the world. Add to this the array of international institutions ranging from the European Union to the more amorphous and decentralized international industry associations and foundations—such as those that govern the Internet—and we have a highly complicated but extremely relevant set of context actors.

Public investment into infrastructure can also matter. For example, historically, investments into the canal system of the United States’ great lakes, along with the Panama Canal, opened up the waterways that enabled steamships to link trade from previously unconnected areas and fundamentally alter the global economy. Or for the case of the automobile, the United States’ investments into creating nationwide highways in the 1950s under the Eisenhower administration, initially to be able to move military supplies around quickly during the Cold War, was the impetus for transforming long distance transport away from railroads to trucks, automobile, and airplanes. (Many parts of US passenger rail are significantly slower than they were 60 or even 100 years ago as the technological implementation choice, shaped by policy, moved away from trains.²³)

Similarly, with something like automated driving, the broad consensus among many researchers is that Level 5, in which a vehicle can navigate to any endpoint completely on its own, regardless of driving conditions, is still in the distant future (some call it the ever receding 20 year horizon since the 1950s), but Level 4, in which there are bounded limitations, such as good weather, the availability of some form of map, and perhaps markings on the road, is easily achievable within 5 years as a technological possibility. **The question is which country or region will invest how**

²³ http://www.slate.com/articles/life/transport/2009/05/stop_this_train.html

much into what form of solution—which may profoundly affect the trajectory of development.

Japan, for example, facing a serious labor shortage driven by its rapidly aging and shrinking population, is likely to politically embrace AI solutions to reduce the number of people needed to perform certain tasks—or even to perform jobs for which there are not enough people. The political opposition expected in regions which have labor surpluses may channel the type of AI applications adopted in these areas.

AI and a Conception of Human Tasks Transformed: the “Algorithmic Revolution”

There is a specific way in which IT tools have been transforming human activity, and AI will accelerate this dynamic.

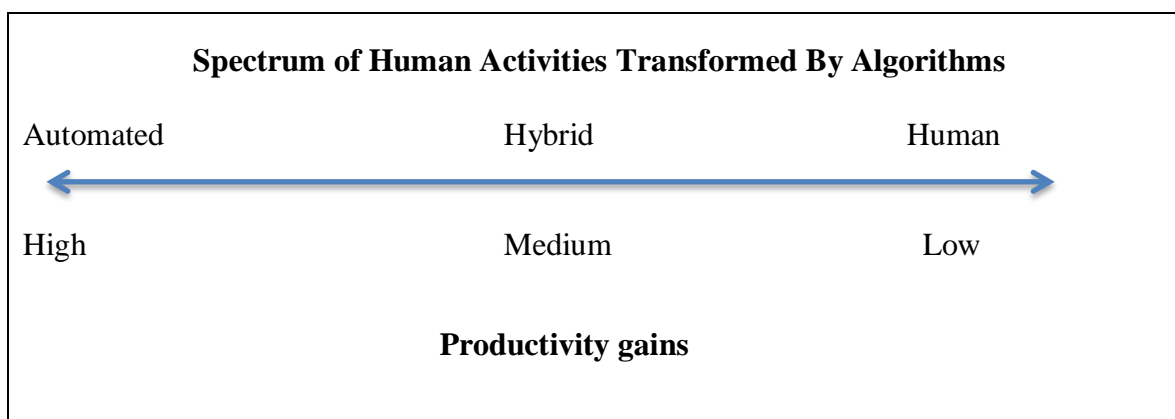
An increasing swath of human activity is captured by algorithms, which allows it to be split apart, transformed, altered, and recombined—the Algorithmic Revolution.²⁴ Human activities can be placed along a spectrum of how they are transformed by algorithms.²⁵

On the one hand are fully automated activities. Data searches, communications, and routine accounting, for example.

On the other hand are human activities that cannot (yet) be fully replaced by algorithms, such as haircuts or replacing carpets in homes. In between are hybrid activities, in which human activities are substantially enhanced by algorithms.

²⁴ Zysman, J. (2006). "The Algorithmic Revolution - The Fourth Service Transformation." Communications of the ACM 49(7): 48.

²⁵ Zysman, John, Stuart Feldman, Kenji E. Kushida, Jonathan Murray, and Niels Christian Nielsen. "Services with Everything: The ICT-Enabled Digital Transformation of Services." Chapter 4 in The Third Globalization: Can Wealthy Nations Stay Rich in the Twenty-First Century? edited by Dan Breznitz and John Zysman. 99-129. New York, NY: Oxford University Press, 2013. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1863550



The range of activities that are moving from human to hybrid and automated activities is growing quickly as the Algorithmic Revolution proceeds with ever-increasing availability of processing power and storage.

AI radically accelerates the Algorithmic Revolution by allowing human activity to be captured far more easily than ever before. Rather than creating a deep understanding of how human activities are performed, AI can find correlations and capture activities. Jeff Bezos, CEO of Amazon, put this eloquently in his letter to shareholders in 2017.

"Over the past decades computers have broadly automated tasks that programmers could describe with clear rules and algorithms. Modern machine learning techniques now allow us to do the same for tasks where describing the precise rules is much harder." (Jeff Bezos 2017)²⁶

It is in this vein that we see activities that could previously only be performed by humans enter into the realm of hybrid, and then automated activities. Uber or Lyft drivers are hybrid in that the skills of learning where to pick up passengers, how to get to their destinations, (and in some countries, how to negotiate for large tips), are automated—although driving from point A to B still takes the same amount of time. By catching times of peak demand, drivers can even increase their productivity by earning more per ride. Eventually, driving will be fully automated.

Even folding clothes, a core human activity that many considered to encapsulate the pinnacle of activities that only humans could perform—William Baumol in his famous paper from the 1960s about how the productivity of service sectors is limited, even used this example²⁷--has recently been the target for automation. Startup companies such as Japan's Seven Dreamers, teaming up with Panasonic, is rolling out a fully automatic clothes folding machine, the Laundroid, that uses AI to learn and recognize garments, using mechanical arms to fold them.²⁸

²⁶ <http://www.businessinsider.com/jeff-bezos-shareholder-letter-on-ai-and-machine-learning-2017-4>

²⁷ Baumol, W. J. (1967). "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis." *American Economic Review* 57: 415-426.

²⁸ <https://laundroid.sevendreamers.com/en/>

AI vs. IA (Intelligence Augmentation), and Jobs

A common concern about AI is how fast it will replace jobs, and what types of jobs will remain. A prominent researcher framed the question as not “which jobs will be replaced” but rather, “which jobs would possibly not be replaced” as a better starting point for inquiry.

Arguably the most cited work for specific numbers of jobs lost is a working paper from Oxford scholars Osbourne and Frey in 2013, predicting that 47% of US jobs will be lost in 20 years (2033).²⁹ The specificity of numbers led to an extensive amount of media coverage, and other reports took this as a starting point. The paper itself has grave methodological problems. As some critics note, it was not published in a peer reviewed publication, which in and of itself may not be a problem, but may suggest methodological weaknesses. In fact, rather than actually looking at the occupational categories in the US (there are 702) and assessing and evaluating each one’s likelihood of replacement, it instead assesses jobs based on requirements of manual dexterity and social perceptiveness; if below a certain score, they predict it will be automated, and if above, they predict it will remain a human activity. As a researcher from the Information Technology and Innovation Foundation notes, this “methodology produces results that make little sense, as when they predict that technologies such as robots will eliminate the jobs of fashion models, manicurists, carpet installers, barbers, and school bus drivers.”³⁰

Another highly cited study is by the McKinsey Global Institute. Many interpret the results as saying that 45% of jobs will be automated, but the actual report predicts that only less than 5% will be fully automated, with the rest coming from shares of employee time. Therefore, certain portions of doctors’ time will be saved by automation, which will amount to the equivalent of a certain number of doctors’ full time work, but it does not predict that this number of doctors will be displaced.³¹

From the early days of AI development, there have been two contrasting approaches to automating human activity. First is the traditional AI, in which people are replaced. Another branch is that of ***Intelligence Augmentation (IA), in which people remain at the core, with their abilities amplified.***³² This has significant implications for how to think about the future of jobs.

For many tasks, IA will enable unskilled workers to perform the tasks of highly skilled workers. Some highly skilled workers will be completely displaced, but others will simply be able to use their time for more productive activities. For example, if medical imaging diagnostics is more effective using AI tools, then surgeons might focus more on performing surgeries and communicating with families.

The question is *how many new opportunities for low-skill workers will be created by IA, and how does that compare with jobs automated by AI?* How will tasks performed by individual workers break apart and be recombined?

Take Komatsu, the Japanese construction equipment company, for example. It currently requires ten years of experience for a worker to master performance of

²⁹https://www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf

³⁰ https://itif.org/publications/2017/08/07/unfortunately-technology-will-not-eliminate-many-jobs?mc_cid=7ec3ca683a&mc_eid=03f4769fd2

³¹ *ibid.*

³² Markoff, John. *Machines of loving grace: The quest for common ground between humans and robots.* HarperCollins Publishers, 2016.

certain cuts into the ground—using the circular motion of a power shovel to dig a slope at a particular angle, for example. However, with sensors embedded throughout the machine, Komatsu can now use workers who are almost completely untrained to perform these tasks. The machine will stop them from making an error, and an autopilot will automate some of the more complex maneuvers. Yet, tasks such as assessing the stability of the ground or identifying whether the large shape near the dig site is a rock or plastic bag can be easily done by an unskilled operator.

Taken as a paradigm, it is fundamentally unknowable at this time how many jobs that are currently high skilled can be performed by low skilled workers,³³ since companies around the world, large and small, are rushing to create such IA systems.

Conclusion and Implications

The purpose of this overview was to provide some information to raise further questions and spark informed discussions and inquiry. We will end simply with several questions.

For companies: are you preparing for commodity AI tools? Are you prepared for tools such as those used by DeepMind to optimize Google datacenter cooling and apply them to various areas of your organization? Who will identify the tools, and who will make sure they are implemented? Do you have a process for empowering the people in your organization who interact with customers or core areas of the business who are not IT experts? **The next phase of AI is likely to empower people without specialized knowledge of AI itself to ask valuable questions and design solutions.**

For places: what are the technological choices that you are wittingly or unwittingly supporting? For example, there are good reasons for strict privacy laws. The question is how to make these decisions while understanding the potential costs of not being able to attain scale in deploying tools or platforms. What are the policy choices that may shape the trajectories of technological deployment more broadly rather than having to adapt the technology to fit local regulations in a way that makes it more difficult to harness the frontier?

For companies as well as places: are the decisions you are currently making following the logic of computing power abundance, or flying in the face of it? Many companies and governments invest vast sums into IT systems that end up being proprietary, with the countries locked into costly long-term contracts. In an era of processing power abundance, every IT-related decision should withstand the question of whether it will withstand the doubling of processing power every 18 months that will enable AI to perform new analyses using new sources of data.

For everybody: What are portions of tasks that are best automated, and how are employment and labor systems able to harness it (not simply automating people, but amplifying people)? Almost every job has portions of tasks that can be automated, and many skills that are currently high end, high skilled jobs will be able to be performed by lower skilled people amplified by IA tools. As a company, where are you positioned on this? As a country, how do you envision dealing with the transition? As an individual, where do you see yourself positioned on this spectrum? Many of the high end, cutting edge skills of today will be built into the tools of tomorrow. Yet many tasks currently impossible will become realities. What is your vision moving forward?

³³ Zysman, J. and M. Kenney (2014) "Where Will Work Come from in the Era of the Cloud and Big Data?" BRIE Working Paper 2014-6.

Are you falling into the trap of “Weapons of Math Destruction”³⁴ – blind faith in algorithms. This overview did not have the scope to provide an overview of dangers of using probabilistic tools to determine individual outcomes, or the dangers of blind faith in algorithms that may be based on assumptions amplifying inherent biases in the data or deployment. While using AI and algorithms to determine the fate of individuals in areas such as criminal justice and employment may have aggregate benefits, it is critical to understand the assumptions behind the models and be aware that it is the responsibility of society to deal with the errors or biases that result from applying the models.

AI, and the current frontiers of our civilization, will march ahead.

³⁴ O'Neil, C. (2016). Weapons of math destruction: how big data increases inequality and threatens democracy. New York, Crown.