BARINGS
INVESTMENT
INSTITUTE

Artificial Intelligence

TURBO-CHARGING THE NEW DIGITAL WORLD

September 2020

By Colin Moar and Matthew Ward

BARINGS

TABLE OF CONTENTS

EXECUTIVE SUMMARY

TURBO-CHARGING TECHNOLOGICAL DISRPUTION

Building the foundations of Artificial Intelligence systems is a key driver of economic growth today. Specifically, combining Big Data and Machine Learning has the power to transform business models and whole industries. While still early, emerging S-Curves of growth across these industries will justify the very large macroeconomic forecasts of AI's impact. An effective AI strategy may likely soon be considered table stakes for any firm planning for success over the next decade.

FROM MACHINE LEARNING TO MACHINE DOING

While AI and Machine Learning are not new concepts, progress has accelerated rapidly due to the concurrent exponential performance improvements of the computing tools, the growing availability of Big Data sets and the falling costs of data storage. As the models for training and testing algorithms have improved and become more defined, we are now seeing emphasis slowly shift to the deployment of these models in the real world, driving business decisions and competitive advantage.

INGREDIENTS FOR SUCCESS

AI operates in the digital domain, typically with high upfront fixed-costs but very low marginal costs, which means business models built on AI will likely see rapidly accelerating returns as they grow. In other words, bigger is better, so it is important to identify the future winners relatively early if possible.

SUPERIOR ACCESS TO BIGGER & MORE DIVERSE DATASETS

A central part of many business models is the ability to generate profit from end users due to the ownership of high-value proprietary data. If AI can access the same pools of information or expertise, it can not only perform similar functions at a much lower cost, but it can deliver greater insights and recognize patterns that human intelligence cannot.

RISKS FROM REGULATION... & CHINA

Governments and regulators are running hard to catch up with the societal implications of greater collection and sharing of sensitive data to feed the AI engines. The dominance of the large Internet and Cloud platforms also raises difficult anti-trust questions regarding market power and consumer protection. The geopolitical competition between China and the U.S. is in part being fought out through investments in technology, with Artificial Intelligence a critical battleground.

INTRODUCTION

Social, political and economic changes today are all shaped in some way by technology. The origins of these revolutionary transformations reside increasingly in the ways businesses are rebuilding their use of data, tapping into new information sources and developing insights that have until now been impossible to do at a reasonable expense. The process will likely disrupt yesterday's winners and create completely new businesses.

The very nature of the digital world is one in which improvements in performance and efficiency occur at an exponential rate. This is in direct contrast to the physical world, which is asymptotic—and ultimately cannot exceed the physical constraints on growth and economic returns. Disruption is inevitable as the vast majority of existing businesses lack the ability to operate in a rapidly changing world. Artificial Intelligence (AI) tools are the missing pieces for most businesses today, and while many pay lip service to the adoption of AI, very few have both the AI tools and the raw data resources to build a competitive advantage in an increasingly digitalized world.

TURBO-CHARGING TECHNOLOGICAL DISRUPTION

The trends around Big Data, Public Cloud, the Internet of Things and even 5G mobile network investments are well underway and, we have argued, on the cusp of dramatically upending companies across industries where firms learn to harness their power to cut costs and engage customers. The addition of AI, which is developing concurrently, is turbo-charging this transformation. Regardless of which industries they hail from, businesses that actively try to harness AI as part of their competitive strategy, and have the data resources with which to leverage AI, are viewed far and away as among the best places to profit from the accelerating rate of innovation.

AI should be regarded as a new and transformational layer of information technology (IT) infrastructure and a necessity for businesses that want to unlock true benefits of digitizing an enterprise. While the leading Internet companies have been the quickest to adopt AI, every other industry is following suit, and adoption of AI tools will likely soon be seen as table stakes for any company that wants to be able to compete effectively in their industry.

The measurable impacts are expected to be incredibly diverse and range from billions of small efficiency gains to a few headline-grabbing leaps forward, affecting businesses of all sizes every day.

CONSIDER THESE EXAMPLES

Delivering health care

The U.K.'s National Health Service (NHS) has partnered with Sensely to create an app called "Ask NHS," which asks patients about their symptoms and helps locate care services, book doctor appointments and provides access to a library of self-care resources. The app delivers a "double-digit percentage drop in costs," according to the NHS, whose budget for England alone reaches £115 billion.[1] All of this occurred before the COVID pandemic forced doctors to consult remotely with patients to avoid infections, putting such applications front and center in the provision of care for the sick.

Weeding out bad actors on social media

Twitter identified 300,000 user accounts with terrorist links in the first half of 2017, 95% of which were detected by AI machines[2]. The scale and complexity of weeding out bad actors on social media continues to grow, however, and now requires AI to recognize suspicious behavioral patterns by users, identify illegal content and adapt rules to regional standards on data and censorship. As malicious activity on social media platforms may sway elections around the world, the significance of these advances can only grow.

Rebuilding a lost historical treasure

Iconem uses photographs of historical sites, such as the Umayyad Mosque in Damascus, Syria, to build 3D models of the sites. Manually stitching together photos took many hours, but applying automated AI algorithms can digitally rebuild a world destroyed by conflict that would otherwise be lost to future generations.[3]

Transforming Chinese finance

The inertia built into China's traditional banking industry could not keep up with the eCommerce revolution that was rooted in China's rapid adoption of the Internet. The digitizing of transactions naturally lends itself to AI tools for detecting fraud on a massive scale, personalized product recommendations on eCommerce sites, and so on. Using AI to disrupt the banking industry to take control of the user relationship has been a critical element that has enabled eCommerce to grow so rapidly in China; the most dominant platforms, Tencent's WeChat and Ant Financial's AliPay, each boast c.1 billion users.

AI's parallels to the breadth and the speed of adoption of electricity at the start of the last century are entirely warranted. As some manufacturers adopted use of electrical motors to enable mass-production processes, spearheaded by Henry Ford, those businesses built on legacy technologies, such as centralized hydraulic-powered machines, suffered heavy losses and even bankruptcy as they tried to fend off the new, lower-cost producers.

Similar scenarios are playing out today, whereby companies and even industries that have been built on the frictions and inefficiencies caused by poor availability of data, small sample sizes or tight ownership of proprietary data sets are being undermined by the proliferation of new digital data troves. Manufacturing and health care are examples of industries ripe for disruption at the hands of companies that aggressively adopt AI-based technologies. Factory owners who fail to adopt Internet of Things[4] solutions and don't generate the data required to automate their manufacturing will likely be overtaken by those who do. The same holds true with drug companies, which must leverage their stores of clinical data to improve drug discovery. These new competitors offer alternative routes to new, more cost-effective and sustainable ways of doing business.

Correct predictions of AI's impact will be a struggle because the potential and the possibilities lie far beyond human capabilities. Moreover, history would suggest that AI's computing capabilities will accelerate exponentially over time. Most of the applications that will be built around AI have yet to be invented.

A couple of simple observations illustrate just how mind-bogglingly huge the challenge is to cope in a digital world:

- Some guess that 90% of all data was created in the prior two years[5]
- Only 1% of all data has been analyzed[6]
- Broad adoption could boost productivity in 2030 by 14%, or $15 trillion[7]

In other words, this ascending data generation means we have only just scratched the surface of extracting value from that data. Humans are incapable of coping with the scale of the challenge, and Artificial Intelligence provides the keys to unlock that value.

A BRIEF HISTORY OF ARTIFICIAL INTELLIGENCE

If all this sounds a little breathless, it's important to review just how quickly humans have arrived at our current AI capabilities. The key innovations in the realm of data science that have sparked the huge resurgence of interest in AI in the last decade came from research into the sub-field of Machine Learning (ML). This is a broad church of different approaches that seek to find relationships, correlations and/or patterns in data. While not a new approach in and of itself, progress has accelerated rapidly due to the concurrent exponential improvements in the performance of the computing tools, the growing availability of Big Data sets and the falling costs of data storage.[8]

FROM MACHINE LEARNING TO MACHINE DOING

AI researchers in the last few decades have split their efforts between building systems where the "rules" for every situation encountered are predetermined, and those where the relationships in the observed data are "learned" and then used to define the machine's reaction when new data or circumstances arrive.

The rules-based approach, however, has largely turned out to be a dead end, as most problems are too large and complex to fit into any pre-defined relationships. For example, the rules an autonomous vehicle must follow to drive in every day, sunny Phoenix are very different from those in Delhi, with its chaotic highways and torrential rain storms.

A breakthrough in machine learning to address image recognition tilted efforts toward a focus on relationships. In 2009, a database called ImageNet was created[9] with millions of images organized and labelled to train and then test algorithms used for image recognition. An annual competition invited researchers to showcase their algorithms built from the ImageNet database.

FIGURE 1: ImageNet Image Recognition

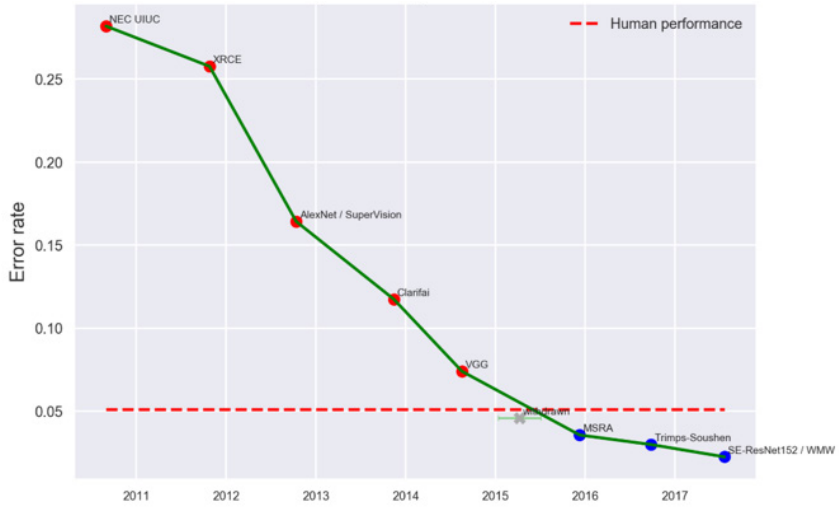

SOURCE: Electronic Frontier Foundation.

In 2012, two graduate students from the University of Toronto, Alex Krizhevsky and Ilya Sutskever, used an approach called a Deep Neural Network to train the algorithm to recognize images (the "AlexNet" point on the chart above). They smashed the competition that year.[10]

In the last year, Machine Learning has embraced yet another notable step up in complexity. Image recognition is relatively simple when compared to the problem of understanding and processing language. While the underlying approach of training an AI algorithm with huge amounts of examples is essentially the same, the number of different parameters the algorithm must learn to recognize in language far exceeds those present in images, and therefore requires algorithmic models that are better suited to the challenge.

Google released a new model called Bidirectional Encoder Representations from Transformers, or BERT, in 2018 to better understand the context of any query entered into its search engine, improving the speed and quality of the results. Now BERT is even being adopted by competitors, including Microsoft. The orders of magnitude, greater complexity in the model parameters, and the data being processed means that the computational power, and therefore the quantity of semiconductors required to power those computers, are much greater than before.

In effect, a neural network is "trained" to correctly identify future images of a dog by being presented with many images of dogs—the more the better.[11] The use of a hierarchy of layers is essentially how human brains learn to identify patterns, as shown in the chart below. Like the brain, a computer starts with simple elements of an image and adds complexity and detail at each layer, always narrowing down the likelihood that it has seen the image before, in this case a dog. The more images are fed to the computer, the faster it gets, which emphasizes the importance of a very large dataset with which to train the computer.

FIGURE 2: How Neural Networks Recognize a Dog in a Photo

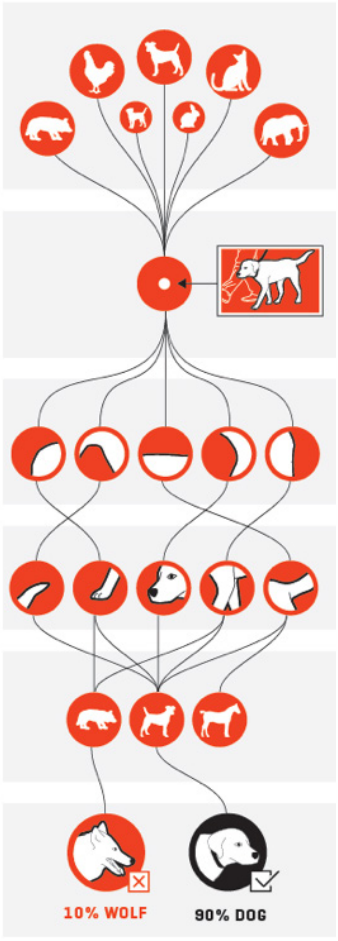

TRAINING
During the training phase, a neutral network is fed thousands of labeled images of various animals, learning to classify them.

INPUT
An unlabeled image is shown to the pretrained network.

FIRST LAYER
The neurons respond to different simple shapes, like edges.

HIGHER LAYER
Neurons respond to more complex structures.

TOP LAYER
Neurons respond to highly complex, abstract concepts that we would identify as different animals.

OUTPUT
The network predicts what the object most likely is, based on its training.

SOURCE: Fortune.

In 2014, two years after AlexNet, a team at Facebook led by Yann LeCun, a leader within the AI industry, built a facial-recognition engine (DeepFace) based on Neural Networks that nearly matched human performance (97.35% accuracy versus 97.5% for humans). A year later, teams from Microsoft and then Google[12] managed to get the error rate for image recognition below that of humans.[13]

Once the Machine Learning process is complete and a model that best describes the relationships among the data points has been defined, the next step is to take that model out into the world and allow it to make a decision based on the new data the machine receives. This process, whereby the model that has been trained and optimized on existing data (Machine Learning) has to then be applied to new data received in the real world, is called "Inference." While still a nascent area in terms of market size,[14] Inference results will begin driving business decisions and likely generate increased revenues and cost savings.

For example, Google's Waymo, which is aggressively pursuing the dream of driverless cars and taxis, has spent the last few years training its AI systems how to drive, using real-world experience recorded through cameras and radar, plus simulated conditions in virtual environments. In December 2018, Waymo deployed their trained AI models in the real world by launching its commercial taxi service in Phoenix. While still employing drivers to supervise the taxi rides, the taxis actually drive themselves as the on-board AI system "infers" what action to take based on real-time data the car senses from its environment. For example, the lights turning green at a crossroads means the car can accelerate forward.

The technology infrastructure required for the training process is very different from what is required to process and make decisions based on live data from the sensors on the taxi. The former needs to cope with billions of data points collected over the decade that constitutes Waymo's database of driving experience. The most popular systems use graphics processors (GPUs), which are inherently suited to large sets of data that can be processed in parallel at huge scale in data centers.

The latter Inference stage, however, has very different requirements. What some call the "Machine Doing" needs to happen at or close to where the decision is applied, as many of the uses will occur in real time. For example, a self-driving car approaching a hazard cannot wait for the sensor data that spotted the neighbor's dog running out in front of the car to be sent over the wireless network, and then out to a distant data center where a computer decides to apply the car brakes. By the time the decision has been relayed back to the car, you are most likely already apologizing to your neighbor for running over their pet. This is where the Inference becomes crucial, so the car can essentially make its own decision based on what it has already "learned" about dogs. This has also helped dramatically lower the barriers to analyzing heterogeneous datasets that were previously trapped in proprietary systems or data silos that could not communicate with each other and limited the usefulness of the data.

To give a sense of the cost dynamic at play here, an on-premise server owned by a single corporate owner will on average be used 15% of the time. In a Public Cloud world where infrastructure is shared across numerous customers, the utilization is on average 65%[15] so you need only pay for one-quarter of the server capacity to do the same work... plus a bit more to cover the profit margin for the Cloud provider.

This dramatic change has launched a whole new model for delivery of software applications. Software-as-a-Service (SaaS) companies, like Salesforce.com, have grown rapidly to optimize this new low-cost infrastructure, accelerating innovation, creating applications to replace older less-efficient solutions, and enable new uses. As these businesses have grown on top of Cloud infrastructure, most if not all of the SaaS companies are incorporating some form of an AI approach into the applications they sell.

Almost as important as falling costs, however, SaaS has undermined the traditional business models of giants like SAP and Oracle, renowned for their closed ecosystems that made it hard to share data with other, unrelated applications.

By its very nature, however, open-source software benefits from creating its own open ecosystem, with many independent software engineers contributing code to the solution. In order to work with existing applications, even those in closed systems, flexibility was needed to integrate with other applications, as well as an ability to cope with many different database management systems and software operating systems, thus effectively removing many of the barriers to adoption for a customer.[16]

This newfound flexibility means that SaaS companies can help their customers combine data from different business silos and free-up the ability to gain valuable insights that were previously hidden, simply because IT systems lacked the ability to talk to one another.

It's hard to underplay just how transformative it is to have this new ability to unlock insights that can be gleaned from data sets that, in the past, were not able to speak to each other, even when owned by the same company.

BIG DATA + INTERNET OF THINGS = INDUSTRY 4.0

The maturation of the concept of an Internet of Things (IoT) has intersected with a point in time where the cost of the sensors and their required connectivity and data transmission have become so affordable that IoT deployments are accelerating rapidly. Where before factories were monitored manually using analogue gauges and warning lights, today we see modern facilities becoming digitalized through deployment of IoT networks, enabling a far greater degree of automation and efficiency. The return on such investments in technology can be significant.

A recent GE survey of 450 field service and IT decision makers found that 70% of companies are not aware when equipment is due for an upgrade or maintenance, and that unplanned downtime can cost companies $250,000/hour.[17]

Tata Consultancy claimed that combining its AI expertise with IoT sensors and digital twin technology (modelling a physical machine digitally, then simulating wear-and-tear over time to predict failures and maintenance schedules) enabled the firm to save a power plant 2–3% on operational and maintenance costs.[18]

To fully unlock the benefits of IoT, the data generated is then delivered to an IoT management application (hosted on-site or, more likely, the Public Cloud). The environmental conditions around the operational areas are ideally AI-enabled to better monitor machine performance and alert the operations team when problems occur, potentially with an automated response to rectify the situation before it becomes catastrophic. As a result, a sharp increase in deployments of IoT solutions are improving the managerial efficiency of manufacturing plants, oil rigs, commercial buildings, and even data centers themselves, which in turn is driving demand for AI tools.

The role of the Public Cloud in the proliferation of IoT is becoming increasingly clear. Both Azure and AWS are actively promoting their computing infrastructure as a way to store, analyze and act upon the data generated by the IoT devices deployed in the real world, or at the "Edge" in IT parlance. Semiconductor companies and Industrial Software companies are investing in the physical, networked sensors and the software tools needed to manage the IoT estate in a factory, or on an oil rig, and leveraging the AI tools offered by Public Cloud vendors. Industrial sectors, which typically spend less than 3% of sales on IT, will see that ratio increase as they adopt AI to increase profits from automated production.

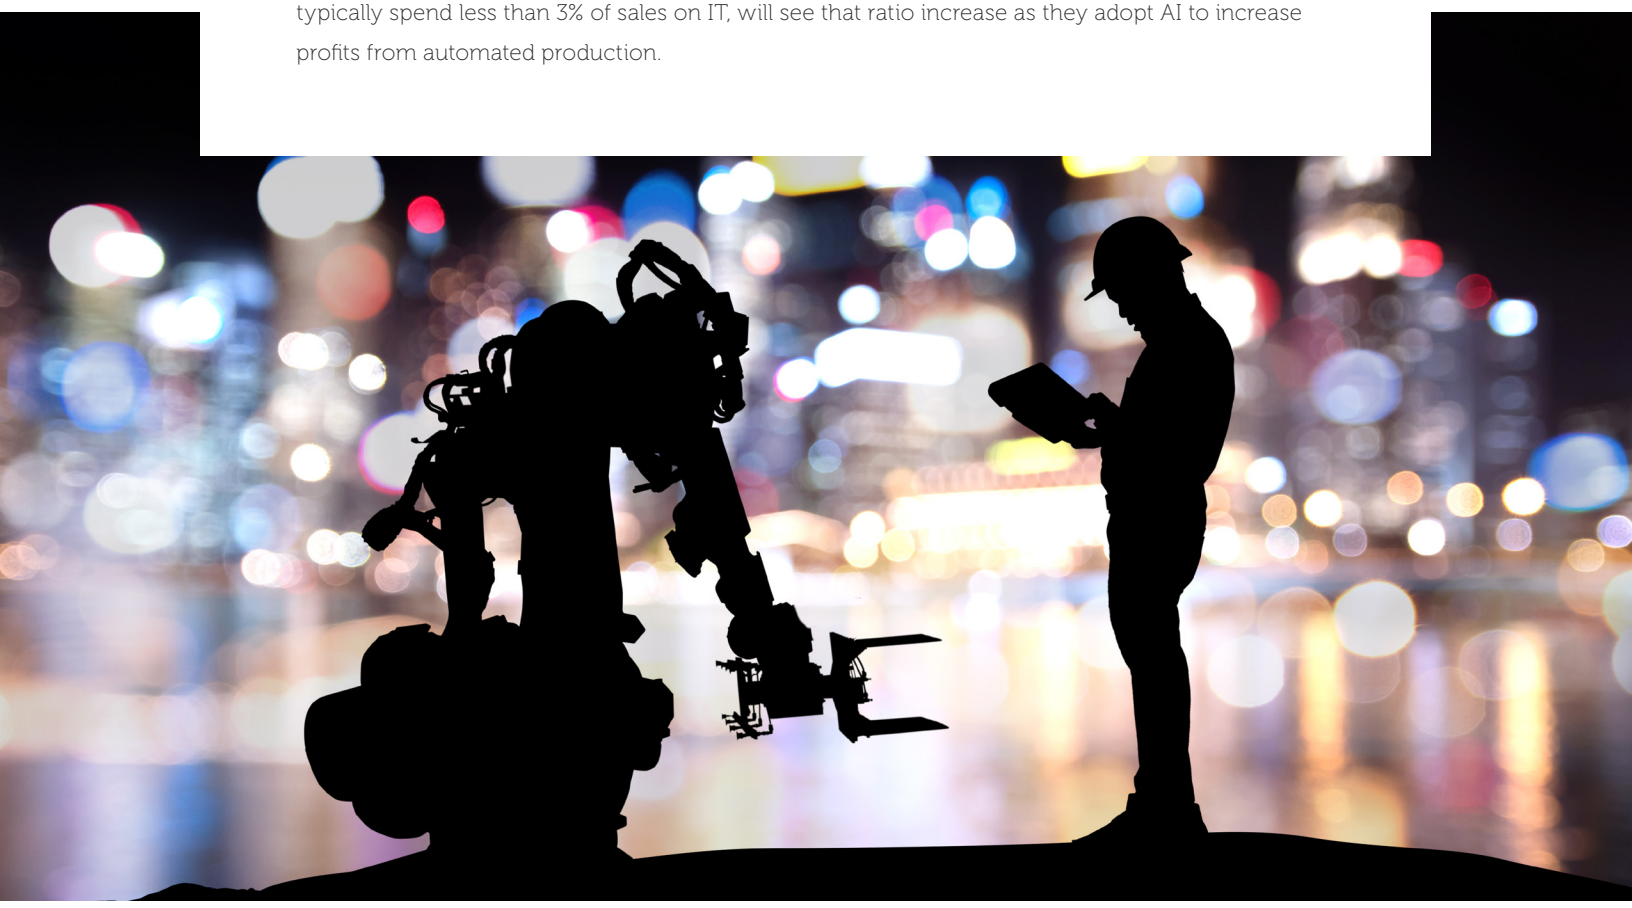

INGREDIENTS FOR SUCCESS

The ease and low cost of adoption are among the most important elements in the recent emergence of AI. While the IT infrastructure required to cope with the volume of data and associated processing power is immense, Public Cloud companies have moved fast to build these capabilities on a grand scale. The scalability of infrastructure unit costs is lower and significantly more affordable than companies that build their own AI infrastructure, allowing even small businesses to afford AI applications to business problems or augment services they can offer to their customers. This is proving to be a powerful force behind the rapid rate of innovation, as companies of all sizes and across all industries are able to access AI and increase their competitive advantage.

THE RIGHT INFRASTRUCTURE FOR THE JOB

The underlying digital data and the software that processes the data reside on top of semiconductors. The falling data storage costs and increased processing power of the servers which analyze the data stems from the way semiconductors improve over time.[19] As Machine Learning has matured and Inference becomes increasingly significant, the choice of which type of semiconductor to use for each task is becoming ever more important.

Intel dominates the market for the central processing units (CPU) used in servers, still holding 90%+ market share. However, as Machine Learning has changed the need to process more data in shorter periods of time, small processors that operate in a parallel manner (CPUs are much larger and operate in a serial fashion) have taken a significant share. NVIDIA's graphics processors (GPUs) are now the leader in parallel computing as, importantly, they augment the power of their silicon chips with a software ecosystem that originally grew out of programming GPUs for computer games. They now count 1.8 million engineers[20] among their CUDA programming framework who could apply their skills to solving AI tasks.

As Machine Learning has matured and the Inference process is starting to emerge as the next key growth area, new semiconductors that are optimized for the different challenges in that setting are being developed. Solutions such as Xilinx's field programmable gate arrays (FPGA) and Intel's Nervana business are competing here. NVIDIA has an advantage due to their incumbency in Machine Learning, but, more importantly, the ability for so many engineers to use their CUDA framework as a foundation for designing the best Inference solutions.

The opportunity in AI semiconductors is huge, especially when the related IoT segment is layered on. Gartner, a technology research and consultancy firm, estimates the market for AI-related semiconductors to grow to $34.3 billion by 2023, from just $4.3 billion in 2018.

An important factor in understanding the breadth of opportunity AI presents is the way in which access to such high-performance computing power has changed in the last decade. All of the major Public Cloud providers now offer customers looking to carry out Machine Learning tasks services built around NVIDIA's GPUs and a number of other AI-focused compute engines. The ability to rent capacity as needed lowers the upfront cost for smaller companies that want to use AI techniques and is a key catalyst for the exploding growth in Machine Learning.

SUPERIOR ACCESS TO BIGGER & MORE DIVERSE DATASETS

Data is the "fuel" on which AI machines run, and without it there would be no AI discussion. The features that define whether a data set qualifies as "Big Data" or not include the volume of the data; velocity of the data (how frequently the data is updated); veracity of the data (is it trustworthy?) and variety (heterogeneity of the different data points in the set). The concept of Big Data therefore captures the idea that there are a rapidly growing number of data sets that exceed a human's ability to analyze it and draw useful insights. A computer is required, preferably a really big and fast one, as are the relevant Machine Learning algorithms to make sense of it all.

These data sets are usually made up of "unstructured data"—in other words, heterogeneous data points which lack the convenient and orderly tagging conventions or predefined data models of the sort of structured data sets that one would find in, for example, an Excel spreadsheet tracking the number and price of widgets sold to particular customers.

Social media's popularity and connection to so much new digital meta-data (the locations of photos that are uploaded, numbers of website "likes," sharing of posts, how long someone spends viewing a given advertisement on their newsfeed, etc.) has led to the development of a whole new family of technologies that are better able to cope with these unstructured data sets. While Oracle dominated the software market for traditional relational database management systems (structured data), new companies such as MongoDB are seeing rapid growth for the databases they build to cope with unstructured data. These databases are a crucial component of any application looking to leverage their stores of data by using AI tools.

The ability to collect, analyze and then monetize the insights pulled from such data points are the engines of profit underneath Facebook and Twitter. As you extend out the sources of data to web search, digital gaming, and to eCommerce in all of its forms (online purchases of travel, food delivery, taxis, books, groceries and so on) you then begin to see a world evolving where digital infrastructure underpins a significant proportion of our society's activity, where the data flowing across the infrastructure justifies how valuable that infrastructure actually is.

For many businesses, a central part of their business model stands astride unavailable or unanalyzed information:

- A stock broker aggregates exclusive data on the capital holdings and flows of their clients to build intelligence on which investments are the most popular
- An advertising agency has proprietary research into consumer behaviors and tastes that drive their ability to offer superior marketing campaigns to their clients
- A drug company has exclusive studies on molecules and how they affect certain diseases, giving them an advantage in finding cures for similar afflictions that emerge in the future

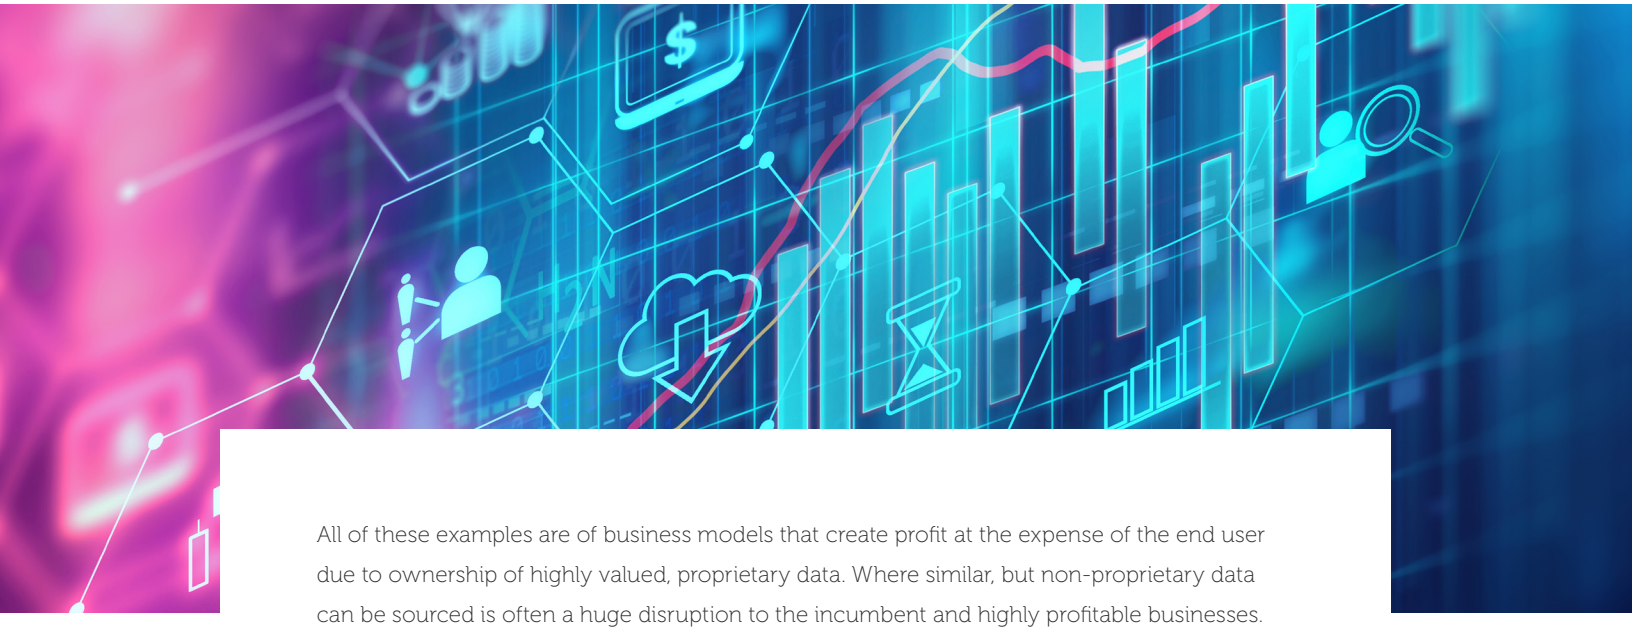

All of these examples are of business models that create profit at the expense of the end user due to ownership of highly valued, proprietary data. Where similar, but non-proprietary data can be sourced is often a huge disruption to the incumbent and highly profitable businesses. The move to electronic trading and wider dissemination of asset prices and other related data has been a key driver of the falling profitability and, ultimately, the consolidation of stockbroking businesses. Even the small-time retail investor benefits from dramatically lower trading costs. The advertising and marketing industries have been thrown into disarray as consumer preference data that used to be sourced through costly and time-consuming (and invasive) market research campaigns is now offered up for free through tracking "like" buttons and measuring click-rates on Internet search results. The sheer volume of data, and therefore the lower statistical error rates on the insights gained from the data, are working to disintermediate the entire section of the industry that used to thrive on the proprietary nature of their hard-won information. A similar story is emerging in pharmaceuticals. New entrants are starting to leverage the availability of electronic medical records and digitized medical journals to identify previously undiscovered insights into how diseases behave, thus accelerating the identification of promising paths for research into new drugs.

If AI can access the same pools of information or expertise, it can not only perform similar functions at a much lower cost, it can also deliver greater insights and recognize unseen patterns that human intelligence cannot.

To justify the effort of applying AI to a problem, there needs to be plenty of data available, preferably where there have also been historical obstacles to combining different but related data sets. This could be due to regulations preventing the sharing of sensitive data, or simply bureaucratic or technological reasons that prevent merging data sets to find new relationships and patterns. Google's recent launch of its Dataset Search tool is a response to this friction and makes nearly 25 million publicly available datasets far more accessible[21] to anyone who cares to make the effort. The outcomes cannot be guessed at, but the opportunity to find new insights into problems is certainly enhanced by Google's efforts to make analysis of data much easier.

HUGE ADDRESSABLE MARKETS

A favorite attribute of technology companies is the playbook of investing a large amount of capital upfront to create a business with near-zero marginal costs. Successful application of AI tools to such a business model should accelerate this process and raise even higher barriers to success for future entrants. Facebook and Google have followed this exact model very successfully.

Large markets are also ripe for the application of AI tools. One of the largest industries already seeing early signs of upheaval is Transport.[22] In 2018, the ride sharing market accounted for just $61 billion of the roughly $3 trillion global transport market that can be addressed by such services, and could reach an estimated $290 billion by 2025, growing at 20% p.a.[23]

The two prior disruptions to the Transport industry—namely cars replacing horses, followed by airplanes replacing long-distance rail and shipping—dramatically reshaped economies as society adapted to the new-use cases enabled by the falling cost and increasing speed of travel and movement.

Today, AI is used by the likes of Uber and Lyft to optimize the matching of demand for riders to the supply of drivers. Tomorrow, AI will also be used to augment the supply of drivers by enabling autonomous vehicles (AVs) to meet the rising demand for transport-as-a-service and improve the economics for rides by removing the need to pay a cut of the ride fare back to a driver.

Crucial to their success, however, is the ability to monetize the data that has been collected and analyzed at great cost and sell it to a vast addressable market.

STRIKING THE OPTIMAL BALANCE BETWEEN HUMAN & MACHINE

AI is not without its limitations. The U.K.'s Government Communications Headquarters (GCHQ)[24] recently commissioned a report from Royal United Services Institute (RUSI) to investigate the role of AI in the U.K.'s National Security.[25] One of the more interesting findings related to the ability of AI to help prevent individual terrorist attacks. The report noted that because these events are relatively rare, their infrequent nature would not enable the AI algorithms to identify robust and repeatable patterns enough to predict future attacks. AI tools applied to pools of Big Data derived from communications monitoring and satellite imagery would however improve the efficiency in collating all of the data required for such analysis. In other words, experts from the intelligence community would be required to work alongside the AI tools and apply their own judgement to better interpret the data put forward through the AI tools.

The flip-side of applying expert supervision is where the task is either very simple but repetitive, dangerous for humans, or involves large volumes of data that go far beyond a human's ability to process (the definition of Big Data).

Shoplifting and other causes of "inventory shrink" cost U.S. retailers $47 billion per year,[26] or 1–2% of sales. Amazon Go's concept is a convenience store with no tills or checkout staff; instead, weight sensors and cameras record what you pick up and link items to your Amazon profile, meaning shoppers can walk straight out of the store with items and be charged later. By using AI with image recognition systems, Amazon is effectively automating the role of the shop security guard to address the 1–2% p.a. impact from shrink that causes so much pain for retailers already operating on single-digit margins.

AI-enabled drone technology is seen as a credible solution to many problems facing farmers with limited resources. Identifying pests, mapping fields' moisture levels, spraying pesticides in a targeted way, and so on are all obvious use cases which require a lot of data and analysis. Using computer-vision systems enhanced with AI tools will help farmers expand their efforts without much incremental investment and therefore help them benefit more from the scale economics of farming. The more targeted use of chemicals and the improved yields from existing farms has obvious environmental benefits.

The automation of repetitive or even dangerous jobs is the source of fears for jobs being lost to machines. In a growing economy with low unemployment, this would actually be positive, as the humans are freed to spend more time on jobs that make the most of their creativity and ability to deal with complexity and uncommon events. The reverse is obviously true in a shrinking economy with significant unemployment. A world in which AI takes a greater role is not necessarily a utopia, but it could be argued that periods of growth can last longer and reach higher than would otherwise be the case.

RISKS FROM REGULATION... & CHINA

As AI embeds itself more into the infrastructure of the economy, it naturally raises important questions around governance and regulation. This is where the questions around AI's impact on employment, data privacy and competitive advantage between economies all become crucial in mapping out the future of this revolutionary technological disruption. Meanwhile, political tensions between China and the United States look set to increasingly shape the development and uses of AI throughout the global economy.

CHINA'S SPUTNIK MOMENT

While the efforts behind AI are global, the two key governance bodies attempting to officiate the AI race are the U.S. and Chinese governments. Their very different cultural starting points and attitudes are leading to stark differences in how AI is proliferating and affecting their societies.

In 1957, America had its Sputnik moment when they realized they were behind the Soviet Union in the space race and had to focus on closing the gap, a push that resulted in putting Neil Armstrong's feet on the moon just 12 years later in 1969. China's burgeoning computer science community had a similar experience when they watched the AlphaGo challenge matches play out.

As Google iterated the AlphaGo approach after the 2016 victory of Lee Sedol, and achieved even greater results with their AlphaZero algorithm and training methods, the world outside of the Silicon Valley computer science community started to pay attention—especially in China, where the State Council in July 2017 released a statement laying out their ambitions for a national effort behind the development of Artificial Intelligence.[27]

U.S. authorities have belatedly noticed China's rapid progress, and it can be argued that many of the recent actions taken under the guise of the trade dispute are an active attempt to stymie China's AI plans. Specifically, the need to have access to the semiconductors used for processing and storing the data is being blocked in some cases where the U.S. cites national security threats from fear of intellectual property (IP) theft. China is significantly weakened in this regard as their domestic semiconductor industry lacks much of the IP required to design and build these chips; the U.S., as the dominant source of supply, thus has a very strong hand in the broader trade negotiations. Hence China's concurrent efforts to eventually become a self-sufficient producer of leading edge semiconductors; China issued a directive aimed at reducing their reliance on U.S. technology, ordering "all government offices and public institutions to remove foreign computer equipment and software within three years."[28]

In China's favor, the sheer scale of users and how much of their lives are conducted on platforms such as Tencent, Alibaba, Meituan-Dianping and the rest, lends itself very well to taking the data points generated by their online activity and applying AI tools to optimize strategies to keep the consumers engaged. All to the profit of the platforms.

One final difference to pay close attention to in what is becoming an AI arms race between the U.S. and China is the two countries' varying attitudes regarding privacy. The U.S. approach to privacy is enshrined in the Constitution, and the missteps by the likes of Facebook and Twitter over protecting users' data have led to Congressional hearings and very public dressing downs of the CEOs. California has followed Europe's lead with more draconian measures to protect consumer data through the California Consumer Privacy Act.[29] The aim here is to give consumers more power over how their data is used, the ability to opt out of data collection by companies and then to set financial penalties should a company be found to transgress.

In contrast, China's progress in AI development is being accelerated by the amount of freedom the platforms have to use the data. For example, Baidu has presented examples of how they monetize the data they collect, including using satellite imagery to track the number of workers going to and from a factory in China that produces products for Apple, then selling the results to hedge funds looking to profit from changing trends in Apple's revenues.[30] The latest example is the coronavirus pandemic and the track-and-trace technology deployed by Chinese authorities to monitor many more aspects of their daily lives than just the contacts with COVID-19 sufferers.[31]

Regardless of one's view of the morality of these decisions, the Chinese are clearly more open to the opportunities to use available data sources, accelerating the improvements in their AI tools and raising the risk that they leap ahead of the U.S. in this critical area of technological development.

ATTEMPTING TO MITIGATE THE RISKS FROM AI

Fake news

The U.S. presidential election in November increases the focus on ways social media can be manipulated to influence the results. Generative Adversarial Networks (GANs) are an AI tool whereby two neural networks try to outdo each other with one generating a fake image, then the "discriminator" tries to identify the fake and provide feedback to the "generator" to help improve the image, iterating the process until a realistic image is created.

GANs are one notable source of "fake news" pictures that have gone viral on social media platforms. The ability to automate this process means the scale of the falsified information generation is enormous, causing deep concern, in particular for the credibility of election results, and therefore a lot of effort is being put into fighting these duplicitous images and videos.[32]

However, GANs can also be a force for good. Training robots how to pick things up by feeding back data on why the past attempts failed can be accelerated, either by using a lot of robot arms at the same time or simulating the process virtually using a GAN.

FIGURE 3: Can you identify which of these images are fake?[33]

The answer is all of the above. Each of these highly realistic images were created by generative adversarial networks, or GANs.

(Note: The bottom right image represents a "class leakage"—where the algorithm possibly confused properties of a dog with a ball—and created a "dogball")

SOURCE: CB Insights.

Training Autonomous Vehicles not to kill people by putting the AVs on public highways is clearly a perilous exercise and one that risks conflict with regulators who will naturally be more risk averse in applying restrictions on how to develop such technologies without harming the public. In March 2018, NVIDIA launched DRIVE Constellation, a photo-realistic simulation system that enables AV systems to be trained using virtual data instead of real-world trials. The ability to train AVs for rare and extreme events or road conditions is an important step forward, especially for those that lack the resources of an Alphabet or Tesla to build large fleets of "learning" vehicles.

An Alphabet-backed start-up, OKWIN, is trying a similar technique for aggregating data from different cancer treatment centers without the patient's data ever leaving the premises. This will inform their efforts at drug discovery in what is obviously a massive global problem, while at the same time addressing concerns around the misuse of sensitive, personal data.

Privacy

As the conversation around data privacy and threats from AI have gained prominence in the media, large technology platforms have been forced to detail how they will approach AI research and application. Europe's recent data protection regime, GDPR, can levy fines up to 4% of revenues for misuse of private data, and regulators elsewhere have called for similar regulations. Alphabet responded with a set of principles to address their AI approach, and are the most detailed,[34] going somewhat beyond "Don't Be Evil." The debate as to where the line needs to be drawn for regulators versus AI companies' responsible self-policing is front-and-center today.

Racial bias

A study published in 2018 showed that Microsoft and IBM's versions of facial recognition software were better at identifying the gender of a white person than of someone with darker skin. Amazon's Rekognition software received similar criticism but has since claimed that improvements to the software removed this "bias" from the results. In September 2019, Amazon CEO Jeff Bezos himself extolled the benefits of facial recognition software while also calling for clear regulations to account for the obvious risk of abuse, which the company is in the process of drafting.

FIGURE 4: Demand for Facial Recognition Tech Rises in China (*Equity and Non-Equity Deals, 2013–2018)*

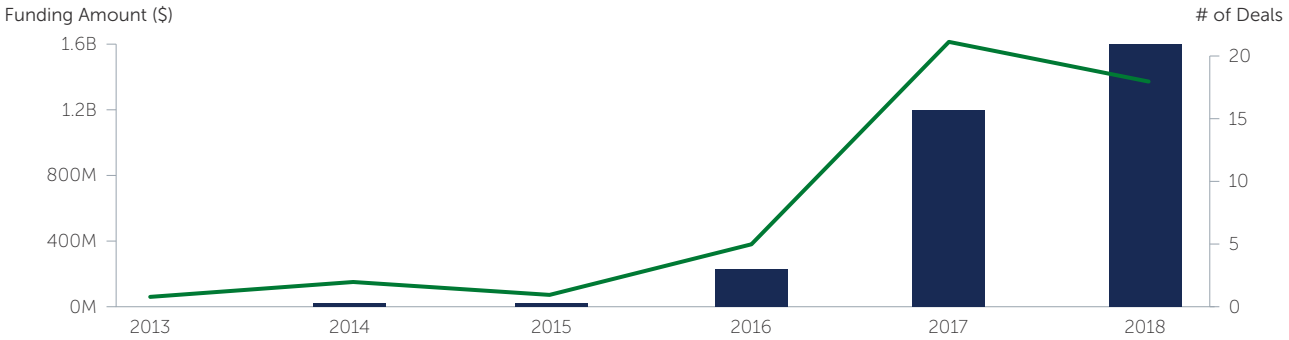

SOURCE: CB Insights.

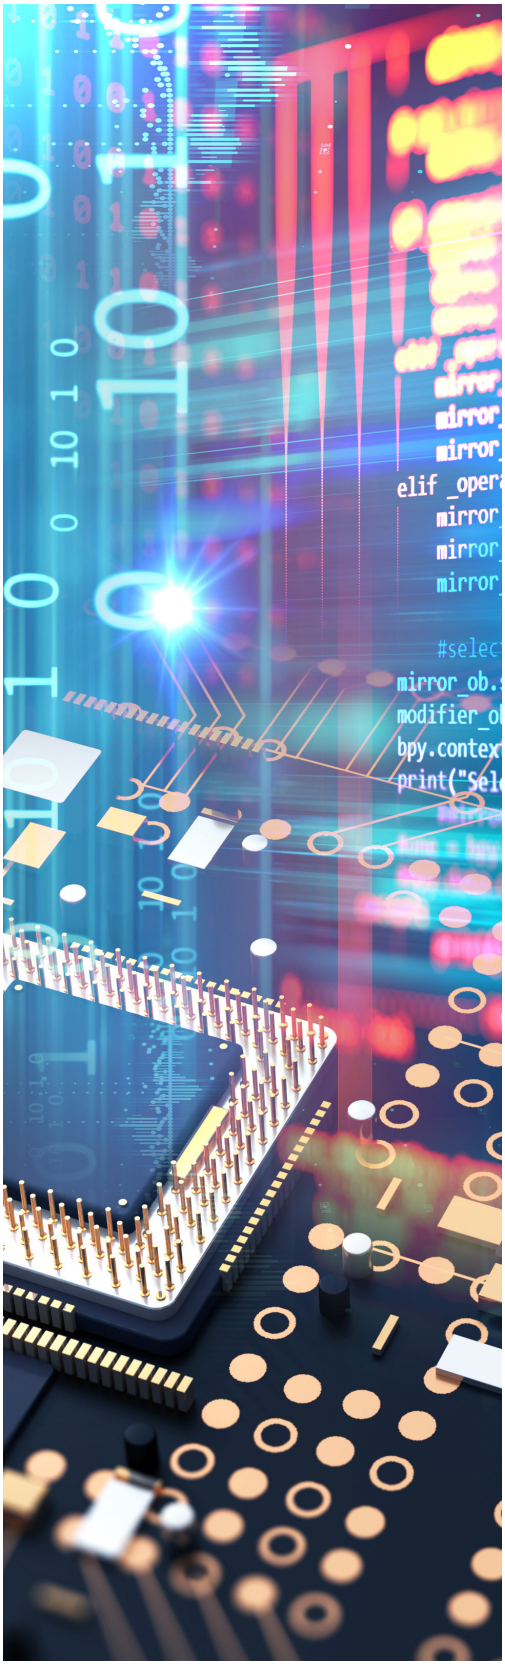

The most recent development here comes in the aftermath of the death of George Floyd and the Black Lives Matter protests. As attention returned to the behavior of the police force and accusations of institutionalized racism, technology companies developing facial recognition software, particularly Amazon, Microsoft and IBM, have decided to stop sales of the technology until clearer rules and procedures are put in place for their use in policing.[35]

There is a powerful tension underlying the responses from the media, regulators and politicians to such fundamental problems with these new technologies. The concept of Moore's Law speaks to the ability of such technologies to improve their performance at exponential rates every year. Those in power cannot possibly keep up with this rate of change, therefore they are forced to act to limit the speed of evolution.

Risks to innovation

As Facebook's CEO Mark Zuckerberg famously noted, clumsy application of regulations simply entrenches the position of the incumbents. The largest players have the experience and the financial resources to build the tools required to meet the new rules.

In contrast, a new entrant sees their upfront fixed cost of entering a market that requires the protection of data rise considerably even before they have generated any revenues, the ultimate effect being to stymie competition and reduce innovation. This is especially pertinent to the AI discussion as the tools required to self-police an Internet platform the size of Facebook or Google (including YouTube) are based on AI. Human monitors cannot possibly cope with the sheer volume of posts that could infringe data regulations.

The lesson from prior technological revolutions is that the regulators and politicians would be caught unprepared as new technologies radically disrupted the existing power balance in an economy. Regulatory responses would inevitably be poorly thought out, rushed and cause damage through unforeseen consequences. The emergence of the Internet and subsequent attempts to set rules around data governance in particular have trodden the very same error-strewn paths.

However, as the coronavirus pandemic worsened it had one positive impact on AI regulation, coming at a time when the EU were seriously contemplating restricting European AI algorithms to only train on European data.[36] Regulators have been forced to acknowledge that the best way to leverage the data available to combat the virus is to apply AI algorithms to all of the available data, without restriction. Many more lives could be saved as a result of a quicker road to finding a vaccine.

CONCLUSION

Artificial Intelligence has exploded since breakthroughs in its ability to perform at super-human levels when asked to identify images. As new and larger data sets are unlocked and access is democratized through increasingly powerful computing resources, it is clear that AI may not only improve the productivity of large swathes of the existing economy but, more excitingly, open up entirely new industries.

While still early in its productive role in society, the opportunity is already there for investors to benefit from seeking those companies which will harness the power of AI and avoiding those set to be disrupted by it. In particular, for many existing businesses, a central part of their business model is their ability to create profit at the expense of the end user due to ownership of highly valued, proprietary data. If AI can get access to the same pools of information or expertise, it can not only perform similar functions at a much lower cost, but it can also deliver greater insights and recognize unseen patterns that human intelligence cannot. These businesses built on proprietary data and closed ecosystems are the ones most at risk from the disruptive impact of AI and therefore the stocks that investors should avoid.

Another reason to focus investors' minds so early in AI's deployments is that digital businesses built upon AI technologies typically suffer high upfront fixed costs but very low marginal costs to supplying their customers. This means the business models built on AI will likely see accelerating returns as they grow. Bigger is better here, so it is important to identify the future winners relatively early if possible.

Similarly, first-movers that have built the stores of Big Data to fuel their AI engines have a significant advantage over start-ups that lack large volumes of relevant training data. This naturally favors the Public Cloud vendors addressing the enterprise market and the Consumer Internet Platforms with billions of users who generate data every day.

This partly explains why governments and regulators are having to run so hard to catch up with the societal implications of greater collection and sharing of sensitive data to feed the AI engines. The dominance of the large Internet and Cloud platforms also raises challenging antitrust questions regarding market power and consumer protection.

With these points in mind, Artificial Intelligence in its many forms is rapidly becoming a part of the underlying infrastructure upon which the broadening digital economy is being built. Macro forecasts for AI's impact over time being measured in trillions of dollars are therefore credible, making Artificial Intelligence an important trend for investing in the coming years and one that should be a key focus for investors today.

MATTHEW WARD

DIRECTOR, HEAD OF TECHNOLOGY SECTOR TEAM

Matthew is Head of the Technology Sector Team and Co-Manager of European Large-Cap Equities. Matthew joined Barings in January 2008 from Bear Stearns International, where he was a Vice President in Equity Research responsible for coverage of Telecoms stocks in the Pan-European region. Matthew holds a BA in Business Economics from University of Exeter, passed the IMC in 2008 and is a Fellow Chartered Accountant (FCA).

COLIN MOAR

ASSOCIATE DIRECTOR, TECHNOLOGY SECTOR TEAM

Colin is a Global Technology Analyst. Colin joined Barings in November 2016 from RWC Partners, where he was an Analyst for a Global Long-Short equity fund. Previously, Colin has held equity analyst and fund management positions at Polar Capital Partners, HSBC Global Asset Management and Aviva Investors. Colin holds a BCom in Business Studies from the University of Edinburgh, passed the IMC in 1997 and the Institute of Investment Management and Research Associate level in 2000.

We would like to take this opportunity to thank Barings colleagues Christopher Smart, Ian Dalton, Agnès Belaisch and Michael Henry for their thoughtful insights in developing the ideas behind this paper and for their helpful contributions through the editorial process.

ENDNOTES

1. Sam Mire, "AI in healthcare Use Case #11: Sensely," Disruptor Daily, June 24, 2019. https://www.disruptordaily.com/ai-in-healthcare-use-case-sensely/
2. Sam Daley, "19 Examples of Artificial Intelligence Shaking Up Business As Usual," Built In, March 25, 2020. https://builtin.com/artificial-intelligence/examples-ai-in-industry
3. Vanessa Ho, "'Heritage activists' preserve global landmarks ruined in war, threatened by time," Microsoft, April 23, 2018. https://news.microsoft.com/transform/heritage-activists-preserve-global-landmarks-ruined-in-war-threatened-by-time/
4. The Internet of Things is where objects and machines have computing devices embedded which enable the measurement and transfer of data.
5. Christo Petrov, "25+ Impressive Big Data Statistics for 2020," techjury, July 19, 2020. https://techjury.net/stats-about/big-data-statistics/#gref
6. Nicolaus Henke, Ari Libarikian, and Bill Wiseman, "Straight talk about big data," McKinsey Quarterly, October 28, 2016. https://www.mckinsey.com/business-functions/mckinsey-digital/our-insights/straight-talk-about-big-data
7. Michael Hoffmann, "Intelligent Economies: AI's transformation of industries and society," The Economist Intelligence Unit, July 29, 2018. https://eiuperspectives.economist.com/technology-innovation/intelligent-economies-ais-transformation-industries-and-society
8. The newfound ability to analyze entire data sets instead of statistical samples can often remove bias and other errors from the output of the analysis. Statisticians could become a rare breed...
9. Wikipedia. 2020. "ImageNet: The ImageNet project was led by AI researcher Fei-Fei Li." Last modified July 13, 2020. https://en.wikipedia.org/wiki/ImageNet#History_of_the_database

ENDNOTES

10. Dave Gershgorn, "The inside story of how AI got good enough to dominate Silicon Valley," Quartz, June 18, 2018. https://qz.com/1307091/the-inside-story-of-how-ai-got-good-enough-to-dominate-silicon-valley/
11. Roger Parloff, "Why Deep Learning Is Suddenly Changing Your Life," Fortune, September 28, 2016. https://fortune.com/longform/ai-artificial-intelligence-deep-machine-learning/
12. Google hired Krizhevsky and Sutskever following their success in the ImageNet competition.
13. Alex Hern, "Computers now better than humans at recognising and sorting images," The Guardian (italicized), May 13, 2015. https://www.theguardian.com/global/2015/may/13/baidu-minwa-supercomputer-better-than-humans-recognising-images
14. For context, at the January 2020 CES show nVidia's CFO noted that Inference was now a double digit percentage of their Datacenter revenue segment.
15. "About AWS: AWS & Sustainability," Amazon, accessed July 8, 2020. https://aws.amazon.com/about-aws/sustainability/
16. This effort was neither easy nor cheap. Software start-ups are often funded by Venture Capital funds that are looking to profit from very high returns that can be generated by successful software companies once the upfront R&D costs of developing a SaaS product have been cleared and the company shifts focus to selling their solution as widely as possible—easy given there are no physical constraints on where they can sell a digital product and the variable costs of selling more units are typically just the sales commission.
17. "After the Fall: The Costs, Causes & Consequences of Unplanned Downtime," Vanson Bourne and ServiceMax, a GE Digital Company, 2017. https://lp.servicemax.com/Vanson-Bourne-Whitepaper-Unplanned-Downtime-LP.html?utm_source=blog&utm_campaign=vansonbourne2017
18. Steve Rogerson, "Tata digital power plant combines AI and IoT," IMC IoT M2M Council, December 11, 2018.https://www.iotm2mcouncil.org/tatadigi
19. This improvement over time is captured by the Moore's Law prediction that transistor counts on a chip will double every 18 months-24 months. A prediction that has held roughly true since 1965.
20. Meeting with nVidia Investor Relations, June 2020.
21. Frederic Lardinois, "Google's Dataset Search comes out of beta," Tech Crunch, January 23, 2020. https://techcrunch.com/2020/01/23/googles-dataset-search-comes-out-of-beta/
22. "On-Demand Transportation Market Size Worth $290.3 Billion by 2025: Grand View Research, Inc.," March 8, 2018. https://www.prnewswire.com/news-releases/on-demand-transportation-market-size-worth-2903-billion-by-2025-grand-view-research-inc-676231383.html
23. "Ride Sharing Market by Type (E-hailing, Station-Based, Car Sharing & Rental), Car Sharing (P2P, Corporate), Service (Navigation, Payment, Information), Micro-Mobility (Bicycle, Scooter), Vehicle Type, and Region - Global Forecast to 2025," Markets and Markets, December 2018. https://www.marketsandmarkets.com/Market-Reports/mobility-on-demand-market-198699113.html
24. Wikipedia. 2020. "GCHQ." Last modified August 5, 2020. https://en.wikipedia.org/wiki/GCHQ
25. Alexander Babuta, Marion Oswald and Ardi Janjeva, "Artificial Intelligence and UK National Security, Policy Considerations," Royal United Services Institute for Defence and Security Studies, April 2020. https://rusi.org/publication/occasional-papers/artificial-intelligence-and-uk-national-security-policy-considerations
26. TJ McCue, "Inventory Shrink Cost The US Retail Industry $46.8 Billion," Fortune, January 21, 2019. https://www.forbes.com/sites/tjmccue/2019/01/31/inventory-shrink-cost-the-us-retail-industry-46-8-billion/#758515216b70
27. "China's New Generation of Artificial Intelligence Development Plan," Foundation for Law & International Affairs, July 30, 2017. https://flia.org/notice-state-council-issuing-new-generation-artificial-intelligence-development-plan/
28. Yuan Yang, Nian Liu, "Beijing orders state offices to replace foreign PCs and software," Financial Times, December 8, 2019. https://www.ft.com/content/b55fc6ee-1787-11ea-8d73-6303645ac406
29. Wikipedia. 2020. "California Consumer Privacy Act." Last modified July 28, 2020. https://en.wikipedia.org/wiki/California_Consumer_Privacy_Act
30. Baidu presented this example at an AI Symposium hosted by Jefferies in London, February 2017.
31. Raymond Zhong, "China's Virus Apps May Outlast the Outbreak, Stirring privacy fears," New York Times, May 26, 2020. https://www.nytimes.com/2020/05/26/technology/china-coronavirus-surveillance.html
32. Facebook's newsfeed has been particularly vulnerable to the spreading of GAN-generated fake images and videos; hence they are pursuing a major effort in partnership with the likes of Microsoft to improve the technology available to detect these. https://ai.facebook.com/blog/deepfake-detection-challenge/
33. "What's Next in AI? Artificial Intelligence Trends," CB Insights, 2019. https://bit.ly/2XyMGr5
34. Sundar Pichai, "AI at Google: our principles," blog.Google, June 7, 2018. https://www.blog.google/technology/ai/ai-principles/; "AI for Good," Microsoft, https://www.microsoft.com/en-us/ai/ai-for-good; "Amazon CEO Jeff Bezos says his company is working on face recognition regulations," CNBC, September 25, 2019. https://www.cnbc.com/2019/09/25/jeff-bezos-says-amazon-is-working-on-face-recognition-regulations.html; AI Facebook. https://research.fb.com/category/facebook-ai-research/
35. "George Floyd: Microsoft bars facial recognition sales to police," BBC News, June 11, 2020. https://www.bbc.co.uk/news/business-53015468
36. Javier Espinoza, "Coronavirus prompts delays and overhaul of EU digital strategy," Financial Times, March 23, 2020. https://www.ft.com/content/f3fd4707-db62-4367-83b1-1b73f1b77862

Barings is a $346+ billion global financial services firm dedicated to meeting the evolving investment and capital needs of our clients and customers. Through active asset management and direct origination, we provide innovative solutions and access to differentiated opportunities across public and private capital markets. A subsidiary of MassMutual, Barings maintains a strong global presence with business and investment professionals located across North America, Europe and Asia Pacific.*

LEARN MORE AT **BARINGS.COM**

**As of June 30, 2020*

20-1296423