AI: AUGMENTATION MORE SO THAN AUTOMATION

A managerial perspective on how firms can effectively deploy human minds and intelligent machines in the workplace.

By Steven M. Miller

The take-up of Artificial Intelligence (AI)-enabled systems in organisations is expanding rapidly. Integrating AI-enabled automation with people into workplace processes and societal systems is a complex and evolving challenge. In this context, we look at management issues related to AI-enabled automation and augmentation, and provide insights relevant to integrating people with smart machines to do cognitive work and other types of work as well.

In this article, we provide an overview of how AI is being deployed inside companies, based on findings from very recently available research on AI applications and managerial impacts. We also highlight adaptation in employee roles as companies increase the usage of AI-enabled systems. This includes state-of-the-art thinking on how humans and machines need to work symbiotically to augment and enhance each other's capabilities. This leads into a discussion on how we fundamentally rethink the standard partnership of human minds and increasingly intelligent machines. We conclude with a statement of the imperative for a new human-machine symbiosis that will help organisations harness the forces associated with the acceleration of change in ways that can lead to growth as well as to productivity increases.

How AI is being deployed inside companies

THREE TYPES OF COGNITIVE AUTOMATION APPLICATIONS

Professor Tom Davenport has spent his professional career doing field studies of new ways in which larger corporations have been making use of information technology and reorganising knowledge work. His book, "Competing on Analytics: The New Science of Winning", originally published in 2007, and revised and republished in 2017, has become a management classic. Over the past four years, he has focused on studying deployments of AI-enabled cognitive systems in the front and back offices of corporations. In a recent joint research with Deloitte, he examined 152 AI deployment projects across multiple industries that are making use of AI-based systems across a wide range of business functions and processes.[1] Based on this work, Davenport categorises AI system applications into three categories:

- **Cognitive Process Automation**: Automation of back office administrative and financial activities using 'Robotic Process Automation'.
- **Cognitive Insights**: Detecting patterns in data and interpreting their meaning using statistically-based machine learning algorithms.
- **Cognitive Engagement**: Engaging employees and/or customers using natural language processing chatbots, intelligent agents and machine learning.

Of the 152 projects, over 84 percent were focused on process automation or analytics while only 16 percent were focused on engagement. Within this smaller set of engagement applications, the majority were for internal employee engagement, not external customer engagement. Davenport notes that as of 2017, most companies were not yet comfortable enough with cognitive systems to 'unleash' cognitive engagement interfaces directly on their customers. My own discussions with companies in Asia that are deploying AI applications are consistent with the findings of Davenport's U.S.-based observations. In the Asian companies I have spoken with, most of the AI applications are for process automation or analytics. For the smaller set of applications focused on engaging with people, I observed the same trend that companies in Asia often start by using AI-based engagement interfaces for supporting their internal employees to test out its capabilities. Davenport cautioned against embarking on highly ambitions moonshot types of AI projects early in a company's experience cycle with AI applications. He stated, "...our study of 152 projects in almost as many companies also reveals that highly ambitious moon shots are less likely to be successful than 'low-hanging fruit' projects that enhance business processes."[2]

INDUSTRY FEEDBACK ON AI USAGE

The 2017 Deloitte State of Cognitive Survey, published in August 2017, was based on interviews with 250 executives and senior managers who are "cognitively aware leaders" within "cognitively active companies" that include "some of the most aggressive adopters of cognitive technologies" in the U.S.[3] Other credible surveys have assessed the degree of AI implementation across industries in Asia to be substantially behind the degree of implementation in North America or Europe.[4] As such, this Deloitte survey report will be very helpful to senior managers and executives in Asian companies considering or already piloting AI applications, or still in the early stage of scaling up deployments across their business. Three of the findings from this survey that I would like to highlight are as follows:

1. The companies in the survey with the most experience and most expertise in AI applications (akin to those in the 'Fast Lane' in a swimming pool) were much more likely to state the primary benefits of using AI as enhancing the ability to create new products and services or to pursue new markets. In contrast, those companies with the least experience and lowest level of sophistication in AI applications (akin to those who are 'Waders' in a swimming pool) were much more likely to state the primary benefit of using AI as reducing headcount through automation. The Deloitte survey summarises this situation by saying, "...while the Fast Lane wants to innovate, the companies in the Wader segment want to automate." Why is this so? The Deloitte survey report postulates, "... it is telling that so many Fast Lane respondents, who have gained stronger returns from cognitive technologies than Waders, see new revenue opportunities as their main benefit. Perhaps this is because their senior leaders understand the potential of cognitive technologies to improve their products and services." In other words, those companies that lack sufficient experience with using AI to enhance products, services and top line revenue growth opportunities will revert to the more conventional assumption that the primary benefit they can realise is bottom line cost reduction through headcount elimination.
2. Almost half of those interviewed said their top challenge with AI applications related to the difficulties of integrating these new efforts with their existing processes and systems. Companies deploying AI systems based on statistical machine learning methods or on expert/rule-based methods more often referred to integration challenges as their top challenge. This is an important reminder that most AI applications need to be integrated into the existing organisation in order to realise business benefits. It highlights that the majority of your IT department—those personnel who are not AI experts—still have a very important role to play in your organisation. You need them to make it possible for the new AI systems to have access to your various data sources. You also need them to integrate the outputs of the AI systems with your existing legacy systems.
3. Most companies interviewed said they have not experienced substantial job losses to date as a result of deploying AI technologies, and they do not predict substantial job losses in the future either. The survey report notes, "In interviews, most companies say that augmentation has so far been much more common than job elimination through automation."

SINGAPORE EXAMPLES

Two examples of companies in Singapore that are effectively deploying AI and machine learning systems are shown in Figures 1 and 2 below. The first example is Vodien Internet Solutions, a small and medium sized enterprise (SME) with annual revenues around US$10 million per year that provides web hosting services and solutions for organisations in Singapore and the ASEAN region. Vodien's total employment is just under 150 people spread across four ASEAN countries. Over 35,000 companies in ASEAN, mostly SMEs, use Vodien's various web-hosting services, resulting in a base of over 210,000 individual users to support. As a SME, Vodien is an important illustrative example because SMEs account for a significant fraction of both national employment and gross domestic product across all ASEAN countries and across the other Asian countries as well.[5] For this reason, it is important for ASEAN and other Asian countries to pay close attention to the extent of progress their SME base is making in adopting AI in ways that lead to business improvement.

As a SME, Vodien does not have much money to burn. While they are an IT hosting services firm with very capable IT technical staff, they cannot afford to have a dedicated corporate AI team, and to-date, they do not have AI specialists on their staff. Their technical staff and executives did their own self-study to get familiar with AI technology and how AI capabilities could be used within the specific setting of their business. Vodien started its AI journey by concentrating on a few critical internal process automation needs together with the internal analytics needs (Figure 1). They were able to do trials and then proceed to production deployments relatively quickly. In addition, Vodien also partners with their external software solution vendors who had already developed AI capabilities. The combination of both internal and external AI capabilities were incorporated into the specific types of software application products that were being used within Vodien. Based on this successful experience with using AI within a few targeted application areas, they are now in the process of exploring how to use AI capabilities to engage and support their internal employees as well as their customers. This pattern of initially focusing on internal process automation and analytics, and then progressing onward to explore how AI can be used for employee and customer engagement, reflects the same pattern reported by Tom Davenport in his U.S. studies of larger companies.

VODIEN INTERNET SOLUTIONS

Users of the System	Type of AI and ML		
	Process Automation	Analytics Based	Customer Engagement
External			Chatbots for Call Centre Support (under consideration)
Internal	Date Centre Optimisation Support Cybersecurity Response Support	Date Centre Optimisation Planning Cybersecurity Data and Pattern Analysis	Intelligent Support for Call Centre Staff (under consideration)

FIGURE 1 *Source: Author's own*

The second example is the Changi Airport Group (CAG), the corporate entity that manages and develops Singapore's highly acclaimed airport. CAG is a large organisation in terms of total employment, total revenue, and breadth and depth of corporate capability. The total Changi Airport ecosystem, including the many vendors and service providers employed on the airport premises, is many times larger than CAG itself, and the CAG team has to orchestrate and oversee the smooth and reliable functioning of this total ecosystem. CAG is a combination of running a company and running a microcosm of an entire city (as the airport is, in essence, a small city in its special way). Figure 2 describes a subset of the AI projects that CAG has been working on. Like many other organisations, getting started with AI applications, CAG started its journey by concentrating on specific task automation, such as limited chatbot implementation and using AI in image recognition application.

About a year ago, CAG launched a strategic initiative targeted on substantially improving their current ability to predict flight arrival times for long-haul inbound international flights (i.e., inbound flights with trip flight durations longer than five hours). The scheduled arrival time is obviously known in advance, but flight arrivals very often deviate from the scheduled arrival time. Even if the flight departs on time from the originating airport, it could arrive late (or early) for a variety of reasons. And even if it is known that a flight will have a delayed departure from its originating airport, it may or may not be able to make up some amount of the lost time. In short, it is much more difficult than most airport users appreciate to know exactly when a long haul inbound international flight will actually arrive at the gate of its destination. And why is the ability to predict gate arrival times so important to CAG? Because so many processes and operations within the airport are triggered by the plane's actual (versus planned) arrival. To appreciate how many different types of resources within the airport ecosystem need to know and respond to the exact arrival time at the gate, just envision all the things that must happen once the plane touches down on the runway and taxis to the gate; the aircraft parks at the gate; passengers disembark and arrive at the immigration queues; all aspects of baggage handling, including passengers retrieving their bags at the baggage claim; and even at the taxi queue. Then of course, there are all the processes related to the preparation for the next aircraft that is scheduled to depart from that same gate.

As straightforward and basic as it may seem (to non-experts in airport operations) to know exactly what time the inbound flight will arrive at the gate, it has not been possible until now to make this type of prediction with the required accuracy, with enough advance lead (at least two hours in advance, otherwise it may be impossible to change plans) and at the scale of operations required (given the large number of flights that are scheduled to arrive every day) to enable substantial improvements in resource planning and optimisation of airport resources. Over the past 12 months, however, CAG has been able to achieve this prediction capability for its long-haul inbound flights. This has been a great achievement for Changi Airport, and a great demonstration of how AI can be used in ways that will lead to substantial improvements in service delivery quality and productivity. CAG is now in the process of working on a number of extensions and enhancements that build on this foundational capability, from better synchronised resource management for all the people and equipment who are triggered by a gate arrival, to finding ways to get the prediction even a few hours earlier to further improve resource planning, and to applying this same type of approach to the large volume of short-haul flights (within a two- or three-hour flight time) that arrive at Changi.

This was a large AI system effort that had multiple parts. There was the enormous data integration effort that involved bringing together existing types of historical data and real-time data that were already available for IT systems usage. It also involved bringing in new types of historical data and real-time data that were either not being used for this type of purpose, or was not previously available. There was the effort to use AI methods to enhance existing algorithms and models that were already in use for this purpose. There was another stream of work to create new models based on new methods and to combine this with and/or compare this against the existing and enhanced versions of the previously used approach for this type of prediction.

As this was a large project involving so many different types of expertise, and CAG wanted to move quickly, they worked with external technology partners who were familiar with AI methods for this type of prediction and with airport domain settings.

CAG's flight time arrival prediction system is a good example of using AI to create new or vastly improved capabilities without putting existing employees out of a job. By improving Changi Airport service delivery while at the same time improving resource efficiency, this new AI application helps to improve the airport's competitive position, and by doing so, makes it more likely that overall airport demand will continue to increase even with competition from other airports. It would mean that employment for airport service support staff would continue to increase. This project is a good example of how humans and intelligent machines can, and need to, work together to bring about major process improvements and more productive allocation of resources. The AI system provides the predictions of arrival times considering many factors. Airport operations planners and supervisors are needed to make the decisions on what to do once they know with a high degree of certainty the time the flight will actually arrive, based on the prediction of the AI system. Based on everything else they are aware of at the airport, the human planners and supervisors work with the systems to devise an adaptive plan for how to deal with the situation depending on how early or late the flight will be. The complete process from runway landing to gate to immigration to baggage and to taxi gets smoother, easier and more efficient. Deployments for service staff and resources can be planned and optimally allocated. Resource flows are better managed. Costs are reduced and customers are happier.

Changi is also in the process of testing and deploying AI-enabled applications for engaging with its internal employees, and for engaging with customers. The third column in Figure 2 shows that CAG is in the process of introducing chatbots for engaging customers and other internal applications for employee self-service support.

CHANGI AIRPORT GROUP

Users of the System	Type of AI and ML		
	Process Automation	Analytics Based	Customer Engagement
External Customers – Visitors to Changi Airport			Chatbots for Changi Airport (starting with Contact Centre)
Internal Staff – CAG employees – Other Airport Staff	Robotic Process Automation for Finance	Flight Arrival Time Prediction	End User Computing Self-Service Support Chatbots for Other Types of Employee Self-Service Support

FIGURE 2 *Source: Author's own*

How employees need to adapt

FIVE WAYS OF STEPPING

In 2016, Davenport and co-author Julia Kirby published a book, "Only Humans Need Apply: Winners and Losers in the Age of Smart Machines", which identified ways in which the deployment of AI systems in front, middle and back office business settings were resulting in work reorganisation and in human job role changes.[6] Davenport and Kirby introduced a model of 'Five Ways of Stepping', providing five alternatives for employees to renegotiate their relationship with smart machines being used to partially or fully automate cognitive work, and for how employees can realign their workplace contributions. Figure 3 below illustrates these five ways of stepping by Davenport and Kirby in an original way, as their explanations of this framework in their earlier Harvard Business Review article and their follow on book did not contain this type of illustration.

Of these five ways that people can 'step' in response to the inevitable increase in the use of AI-enabled smart machines, the one that involves the largest number of people is the 'step in' role. This is where a person who has deep experience with a process or functional area learns to work symbiotically with the smart software application. For example, the person will know the best sources of data to use to train the machine, or they will know when there are special circumstances that apply that invalidate the system's outputs because these special circumstances are not reflected in the training data, or when factors in the environment have changed sufficiently to warrant an update of the training process and model.

The other four types of stepping are also important, but I believe three out of these other four stepping roles will involve fewer numbers of people than those who need to learn how to step in. 'Step forward' refers to those in the organisation who will be the technology designers and engineers, and the software system creators, consultants and implementers. This includes the analytics specialists and the data science professionals who develop, test and deploy the new models that integrate methodologies from statistics, other areas of advanced analytics, and various fields of AI including machine learning. Across user organisations and technology and service providing vendor organisations, there will indeed be a steadily growing demand for those who can step forward in this way. While new jobs are rapidly being created for this type of role, and this will continue, I still believe it will involve far fewer people than those who will need to step in or to substantially change the nature of their work in other ways.

'Step narrowly' refers to those who will keep on doing some type of highly specialised work they currently do, and continue doing it in much the same way as they have been doing it. While in principle that type of work could be done by AI systems, Davenport and Kirby note there will always be highly specialised tasks where the niche and scope of applicability is so narrow that it will not be worth it for the organisation or for the commercial software vendors to spend efforts on automating the task because the expense of doing so outweighs the benefit. By definition, there cannot be large numbers of people who are able to step narrowly in this way.

'Step up' refers to the executives who will make the decisions on when, where and how to invest in AI capabilities across the portfolio of the organisation's business needs. There will only be a relatively small number of top executives, as well as division level executives and unit level heads and supervisors who are the ones who step forward into this type of role.

If one is not stepping in, stepping forward, stepping narrowly, or stepping up with regard to AI, the Davenport and Kirby framework only leaves one other alternative, and that is to 'step aside'. A productive, income-producing way to step aside is to be the person who handles the human-to-human interface that is often so important in the value chain. For example, increasingly, insurance companies are using AI systems to analyse their claim submissions and to make the initial round of judgements on the disposition of the claim. For example, Prudential Insurance in Singapore announced in November 2017 that it could dramatically reduce the time required to evaluate certain types of insurance claim submissions because of new capabilities provided by AI systems it had recently implemented.[8] A capable human has to be able to explain to important existing customers why they are getting a particular decision from the insurance company, and especially when they do not get the decision they were hoping for, why they should continue on with using this company as their insurer. That person could be an employee involved in claims analysis or sales, and who steps aside to handle new types of situations for customer engagement and communication. Similarly, certain segments of financial advisory are already in the process of being supplemented and even supplanted with AI-based systems with growing degrees of knowledge, analytical capability and intelligence. But even if these algorithmic systems make intelligent use of all the data made available to them to make very well-informed financial plan recommendations, some very capable human has to motivate and persuade the client to accept and follow the plan. Another part of this step aside role is to continually follow up with the client to reinforce the need to follow the plan, as well as to see if the client's situation has changed in ways that would require the creation of a new plan. A similar situation is likely to occur in parts of medical consultation, diagnosis and treatment recommendation. There could be a need for a large number of people to step aside in this way, and to focus on the very last step, and perhaps the most crucial step in the value chain, which is engaging with the customer in ways that close the loop and realise value, both for the customer as well as for the service provider.

A less desirable and non-income producing way to end up stepping aside is to be automated out of a job. We will come back to this issue of AI for automation versus for augmentation in a moment.

The Five Ways of Stepping is a simple yet clever and powerful conceptual framework. It helps managers and executives think through the adaptations that employees at all levels and across all functions will need to make as the deployment of AI applications continues onward, and continues to accelerate. Additional opportunities for stepping that these two business thinkers implied but did not explicitly discuss are the key interface roles between the five stepping roles. For example, I believe there will need to be substantial numbers of people working the interfaces between the 'step in' role, and the four other roles. Maybe the same people doing the 'step in' role will handle this. But it is easy to think of many situations where there will be a need for additional people to be that special humanly intelligent and adaptive 'agent' who serves as the interface between those focused on stepping in, and those focused on steeping forward, stepping aside, stepping up, and even stepping narrowly.

DAVENPORT'S "FIVE WAYS OF STEPPING"

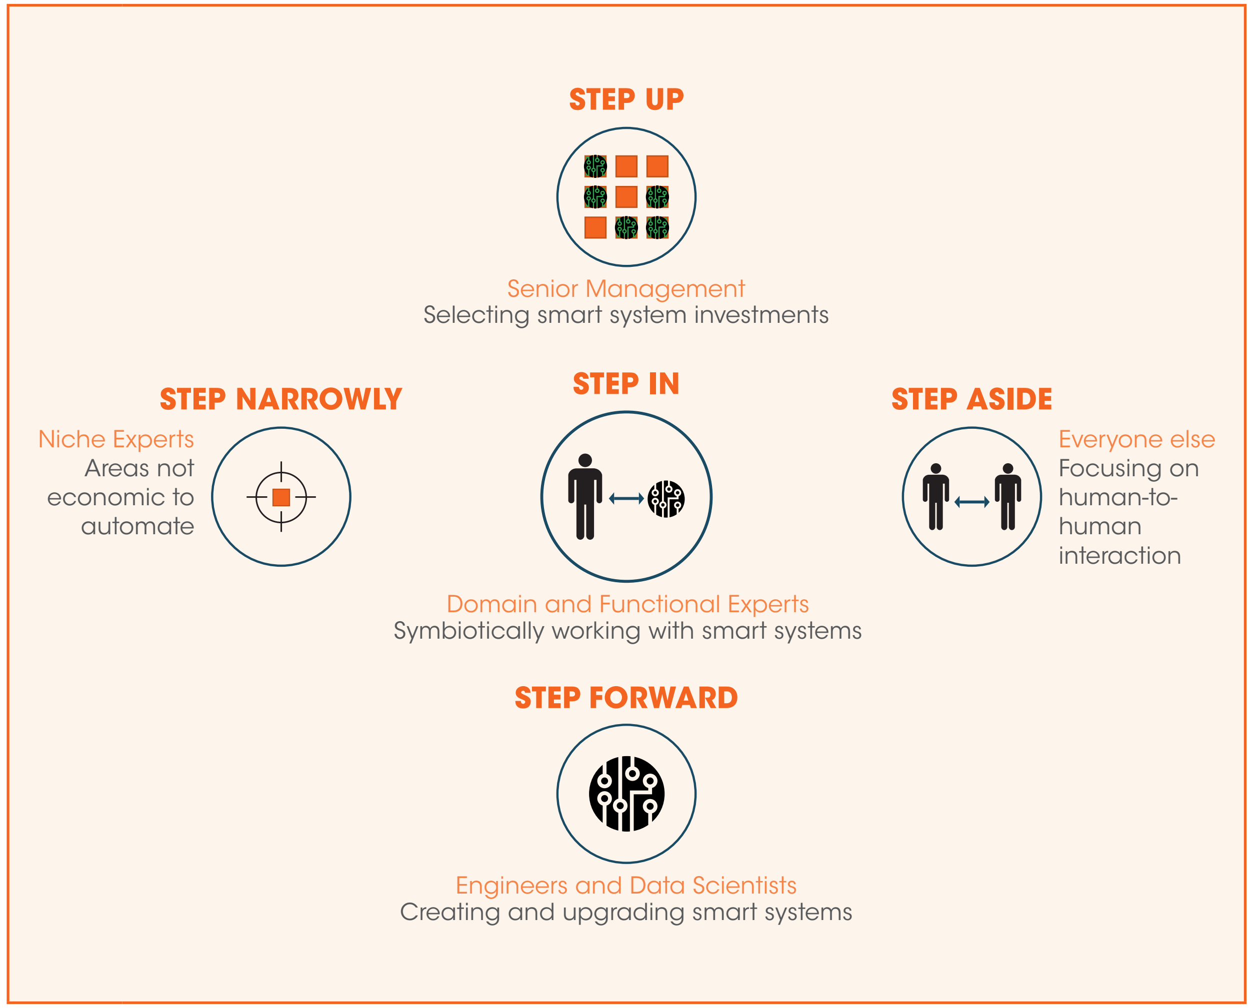

FIGURE 3 *Source: Author's adaptation from the work of Davenport and Kirby (2016)[7]*

AUTOMATION V. AUGMENTATION

As part of the conceptual framework for Five Ways of Stepping, Davenport and Kirby called for an approach to integrate humans and smart machines that emphasised augmentation over automation. As they stated in their June 2015 Harvard Business Review article "Beyond Automation",

> *What if we were to reframe the situation? What if, rather than asking the traditional question: What if tasks currently performed by humans will soon be done more cheaply and rapidly by machines? We ask a new one: What new feats might people achieve if they had better thinking machines to assist them? Instead of seeing work as a zero-sum game with machines taking an even greater share, we might see growing possibilities for employment. We could reframe the threat of automation as an opportunity for augmentation.*[9]

I created two images (Figures 4 and 5) to visually summarise the key points that Davenport and Kirby emphasised regarding an automation mindset versus an augmentation mindset that are important for executives to keep in their mind when they oversee strategic initiatives and investments in AI within their own organisation to improve their processes, services and products.

Davenport and Kirby hit the nail right on the head when they came up with these two summary statements:

- Employees hate automation and love augmentation.
- The same tools can be used to automate or augment, but the intents behind the use of these technologies are 180 degrees apart.

Executives and managers who do not take these two statements to heart will encounter huge obstacles and fierce resistance when they try push AI deployment initiatives forward in their organisations. As noted earlier in this article, from the work of the Economist Intelligence Unit 2017 Global Survey on AI usage in industry, the overall broad base of companies in ASEAN and Asia are still lagging behind their counterparts in North America and Europe with respect to the degree of AI implementation–despite the great success of a relatively small number of well-known AI implementation superpowers in Asia, and particularly in China. As the majority of ASEAN and Asian companies are just setting out on this journey, it is important for them to pay attention to the comments of Davenport and Kirby highlighted above if they want the cooperation and support from the people within their own companies.

Pursuing AI system developments with a mindset of augmentation versus automation makes all the difference to employee morale.

HUMAN AND MACHINE HYBRID ACTIVITY

The concepts of augmented intelligence and human-machine symbiosis can be traced back to the activities of U.S. computing pioneers J. C. R. Licklider and Douglas Englebart in the 1950s, and to their respective writings on these topics that were published in the early 1960s and onward. In 1960, Licklider published an article titled "Man-Computer Symbiosis" where he stated the goal of enabling people and computers to cooperate in making decisions and controlling complex situations

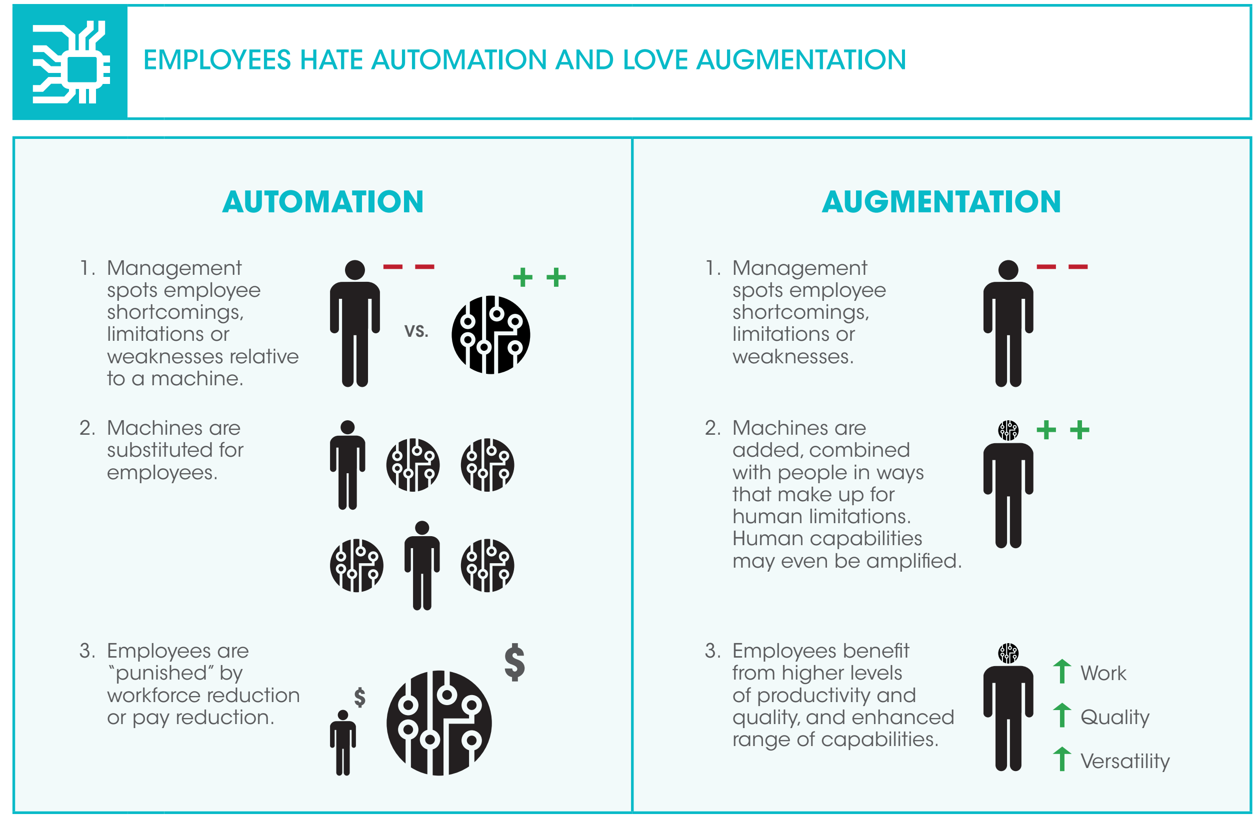

FIGURE 4 — *Source: Author's adaptation from the work of Davenport and Kirby (2016)*

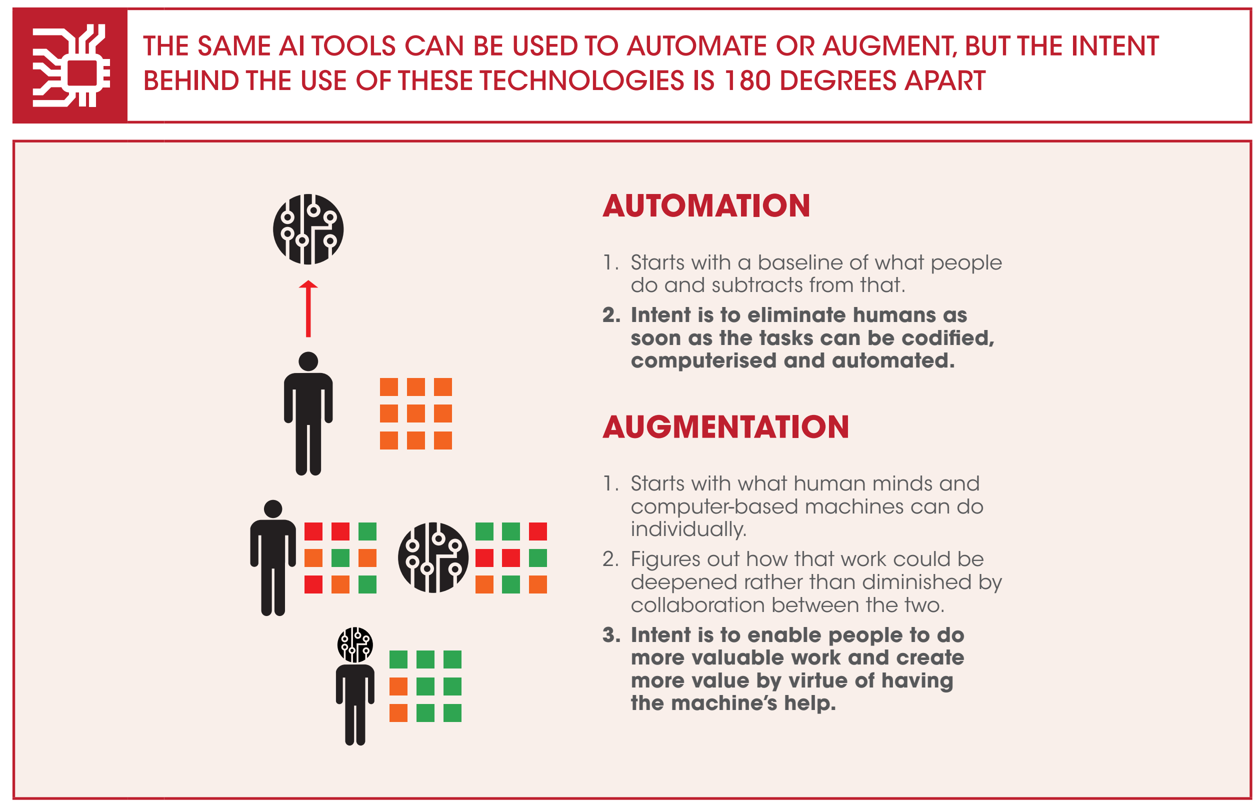

FIGURE 5 — *Source: Author's adaptation from the work of Davenport and Kirby (2016)*

without inflexible dependence on predetermined programmes. Even at that early time, he noted, "Preliminary analyses indicate that the symbiotic partnership will perform intellectual operations much more effectively than man alone can perform them."[10]

Englebart is acknowledged as the creator of the concept of 'Augmented Intelligence.' In 1962 he published a report, "Augmenting Human Intellect: A Conceptual Framework," submitted to the U.S. Air Force Office of Scientific Research, where he described, "...the first phase of a program aimed at developing means to augment the human intellect," based on enabling new types of interactions between people and computers.[11] This report led to the creation of the Augmenting Human Intellect Research Center at Stanford Research Institute in the mid-1960s.[12]

Since the origins of the electronic, programmable computer in the 1940s and 1950s, there have continued to be parallel progressions of work emphasising the use of computing for automation (to replicate and replace human action or thinking) as well as the use of computing for augmentation (to supplement human cognitive capability in a symbiotic manner). Since the 1950s, this same parallel progression has been ongoing with respect to the meaning of AI. There have been researchers and companies focused on demonstrating the use of AI software systems and machines for automating tasks in ways that enable the elimination of human input. There have been other researchers and companies focused on demonstrating the use of Augmented Intelligence for supplementing and amplifying human capabilities for cognition and collaboration. The key point is that the concept of augmenting human intelligence through new types of symbiosis between people and increasingly intelligent machines has a long history going back to the very earliest days of the use of electronic computers.[13]

The Davenport and Kirby work mentioned above emphasises that management should approach the use of AI in their organisations with a mindset of augmentation (symbiotic enhancement) versus that of automation (replacement). But how does a company go about doing that, especially at larger scales of deployment? Is there a framework that can be used to guide this type of effort that is general enough to apply to a wide range of business work settings, but specific enough to provide useful guidance to management? While Davenport and Kirby provide a high level (though very useful) conceptual model for 'Five Ways of Stepping', and they provide a few example descriptions of new ways in which employees are 'stepping in', they do not provide the next level of detail for a supporting framework that management can use to think through and execute a systematic approach to augmentation that can be applied across all parts of the business. Very recently, two senior consultants from Accenture, Paul Daugherty and Jim Wilson, did an excellent job of providing this type of framework with their new book, "Humans + Machine: Reimagining Work in the Age of AI" that was published in March 2018.

After collecting data from 1500 practitioners involved in various types of advanced automation, AI and business transformation efforts, and looking at 450 projects in more detail through observational and case studies, they concluded, "The simple truth is that machines are not taking over the world, nor are they obviating the need for humans in the workplace. In this current era of business process transformation, AI systems are not wholesale replacing us; rather, they are amplifying our skills and collaborating with us to achieve productivity gains that have previously not been possible."[14]

Daugherty and Wilson describe the ongoing emergence of a third wave of business process transformation efforts that are resulting in processes that are more flexible, adaptive and faster than was previously possible, enabled by digitalisation, AI, and other ongoing technology developments.[15] Their key finding is that this ongoing AI-enabled transformation effort, "...has created a huge, dynamic, and diverse space in which humans and machines collaborate to attain orders-of-magnitude increases in business performance. We call this the 'missing middle'– 'missing' because almost no one talks about it, and only a small fraction of companies are working to fill this crucial gap."[16]

They further elaborate, "In the missing middle, humans and machines aren't adversaries, fighting for each other's jobs. Instead, they are symbiotic partners, each pushing the other to higher levels of performance. Moreover, in the missing middle, companies can reimagine their business processes to take advantage of collaborative teams of humans working alongside machines."[17]

In their framework, jobs that require high degrees of leading, empathising, open-ended and/or less structured creating, and holistic judging (as in values-based, strategy-based or context dependent) will remain as human-only activities. On the other side of the spectrum, jobs that require high degrees of transaction execution (especially at large scale), repeatable iteration, model-based prediction (where the necessary digital data is available) and adaptions that can be specified by models or rules or be derived from data, will increasingly become machine-only activities.

Their missing middle is the growing range of work activities that are best done as a hybrid of human-machine effort, in the middle of the spectrum between those activities that remain as human-only and machine-only. The examples in their new book provide the current day as well as emerging future version of the symbiosis concept first articulated by Liklider in 1960, and by Englebart in 1962. Their framework for the missing middle helps us to elaborate on a framework for the 'step in' role first described by Davenport and Kirby that is central to their emphasis on augmentation over automation.

Daugherty and Wilson divide the missing middle into two parts (Figure 6). One part is where humans are augmenting the capabilities of AI-based systems and machines by providing the symbiotic activities of training, explaining or sustaining. For example, humans need to work out the training strategy by which the AI system will be able to initially achieve acceptable levels of task performance, and then continue on to improve upon its level of task performance. The human trainer needs to arrange for the AI system to have access to the data it needs for task performance-related training. There are other dimensions to training as well. The human can provide the system with feedback on how to engage with other people in a way that seems 'more human,' as in, with more empathy, or with more cultural or diversity-related sensitivity.

Humans are needed to help explain the output of an AI system. For various types of internal and external audit and accountability reviews, and for many other reasons, it is often necessary in business and work settings to be able to explain why a recommendation, a decision, or a prediction from an AI system or machine is what it is. Daugherty and Wilson describe additional aspects of how humans need to help with explaining. For example, a human 'explainability' strategist has to think through the transparency of the AI methods used, and the trade-offs between ease of 'explainability' versus model accuracy and related task performance. The human explainer also needs to do 'output autopsies' to understand what happened when certain outputs of the AI system are wrong or off base, or inappropriate. This diagnosis of why a wrong or unacceptable answer occurred has to be fed back to those people responsible for the AI system's training and design.

Humans need to sustain the AI system and the process into which it is embedded by ensuring that the AI is always and only used in proper ways. This means supervising process performance and AI system outputs to make sure there is adherence to compliance requirements and ethical standards, and that overall, the AI is only being used in ways that serve people in their work in appropriate ways.

The second part of the missing middle is where human capabilities are extended or enhanced. The AI system or machine supports the human in a way that makes it appear as if the human has some degree of 'superpower'. In the Daugherty and Wilson framework, this occurs through the ability of the AI system to amplify, to interact, and to physically embody in ways that are beyond the

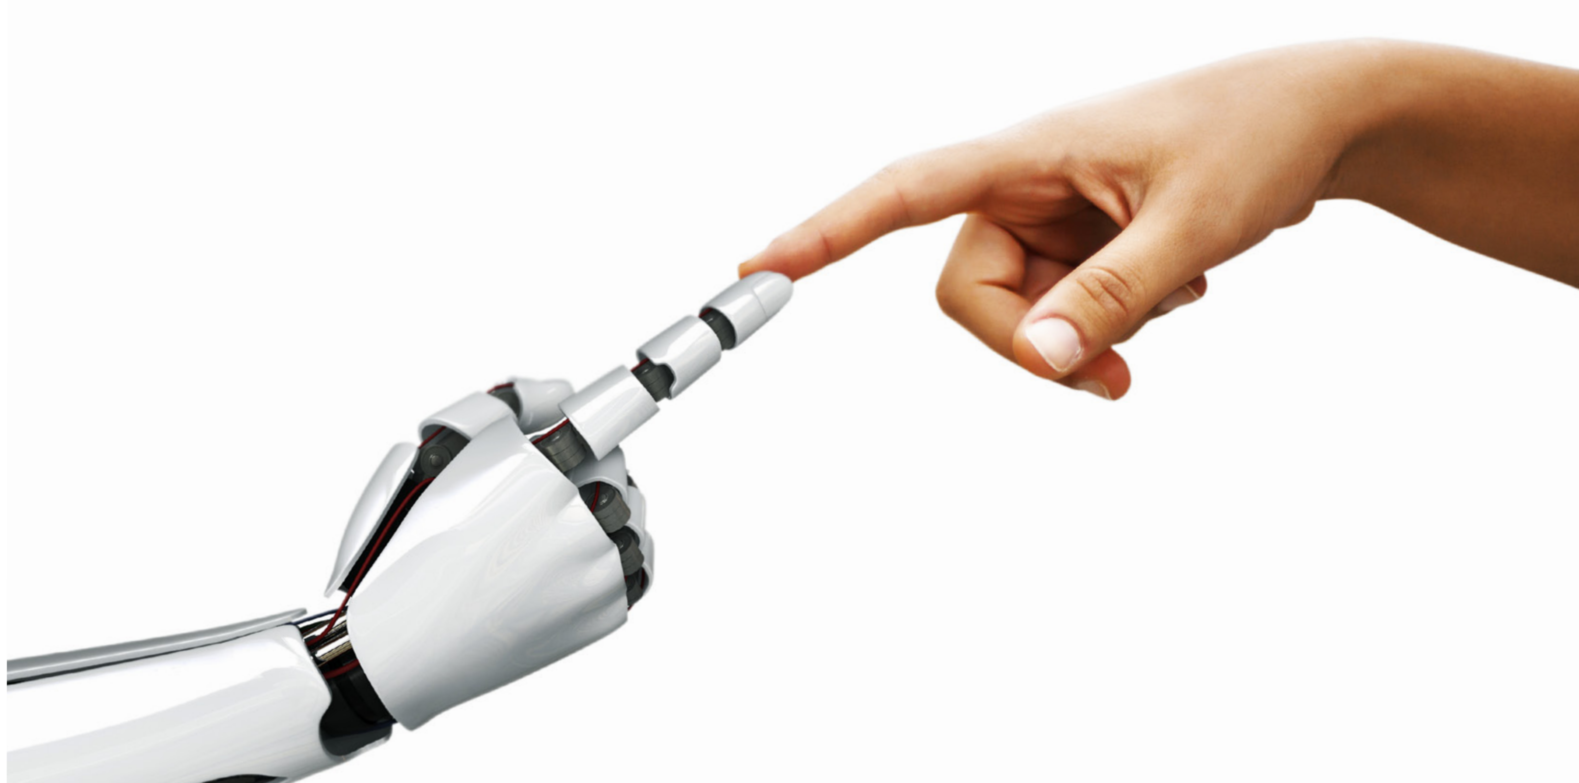

abilities of any human. For example, AI-enabled analytics systems can process vast amounts of data at superfast speed, and use computational methods and models to quickly identify patterns, make predictions, or make choices based on consistently logical analysis of all this information. In this way, the AI system amplifies the cognitive capabilities of a human. Using extended reality (augmented reality and virtual reality), a human can engage with vast amounts of information or with high dimensional information in ways that would not be possible otherwise. Or a person can use gestures or voice to engage with machines, or with people in other locations. This is how AI systems augment people through enhanced ability for interaction. Mobile robots can bring things to people, or can carry heavy loads for people, or help with physical parts of tasks that are damaging to people. These are examples of how AI-enabled machines augment humans through physical embodiment.

Daugherty and Wilson call for humans and AI-enabled systems and machines to fill the missing middle by working closely together in new types of roles with new kinds of collaborative partnerships. They provide a strong incentive for moving in this direction, and additional support for why managers should use AI as a means for augmentation rather than as a means for automation. They conclude, “In our research, we found that companies that use AI to augment their human talent while reimagining their business processes achieve step gains in performance, propelling themselves to the forefront of their industries. Firms that continue deploying AI merely to automate in a traditional fashion may see some performance benefits, but those improvements will eventually stall. We predict that within the next decade, there will emerge a vast difference between the winners and losers, a difference determined not on whether an organisation has implemented AI, but on how it’s done it.”[18]

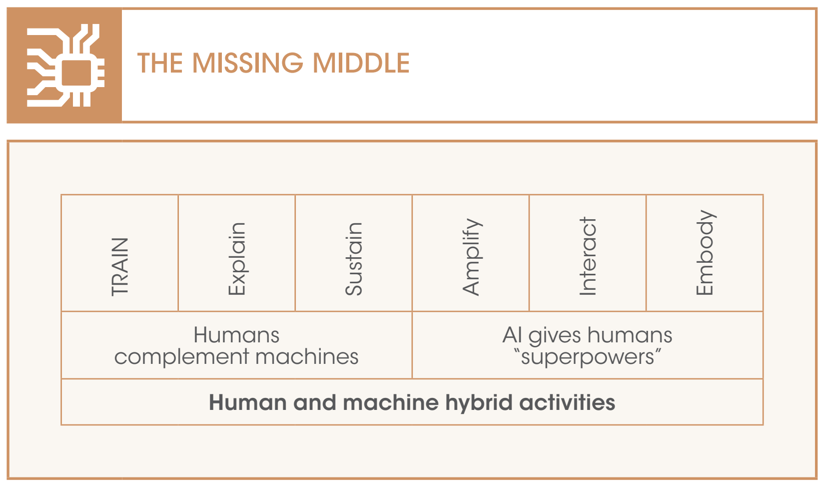

FIGURE 6 *Source: Author’s adaptation from the work of Daugherty and Wilson (2018)*

Rethinking the standard partnership

In 2017, MIT Sloan School researchers, Andrew McAfee and Erik Brynjolfsson, published the book, “Machine, Platform, Crowd: Harnessing our Digital Future,” which was their third book on factors driving the digital economy.[19] It is divided into three parts—Mind and Machine, Product and Platform, and Core and Crowd. Each part describes the great rebalancing that is presently taking place across industries throughout the global economy and more broadly across society. Their premise is that these three concurrent and interrelated rebalancings are the main forces driving current, emerging and foreseeably future trends in the digital economy.

The section on Mind and Machine focuses on the changing relationship between human minds and increasingly intelligent machines in the workplace. McAfee and Brynjolfsson articulate what they call the “standard partnership between minds and machines” that has developed within organisations (in both the private sector and the public sector) that have been using computers to automate work tasks since the 1970s.[20] Their key point is that this standard partnership between minds and machines that has been carried out in our workplaces all over the world over the past 50 years must change and, in fact, is rapidly changing.

The McAfee and Brynjolfsson statement of the standard partnership between human minds and computing machines that has existed until recently is summarised in a visual image (Figure 7). The top part of the image states their observations on the standard partnership. The bottom part of the image states their views on the limitations of this standard partnership.

TWO MODES OF HUMAN THINKING

McAfee and Brynjolfsson refer to System 1 and System 2 modes of human thinking, and use these concepts to explain important limitations of human judgement and decision-making vis-à-vis data driven, algorithmic and

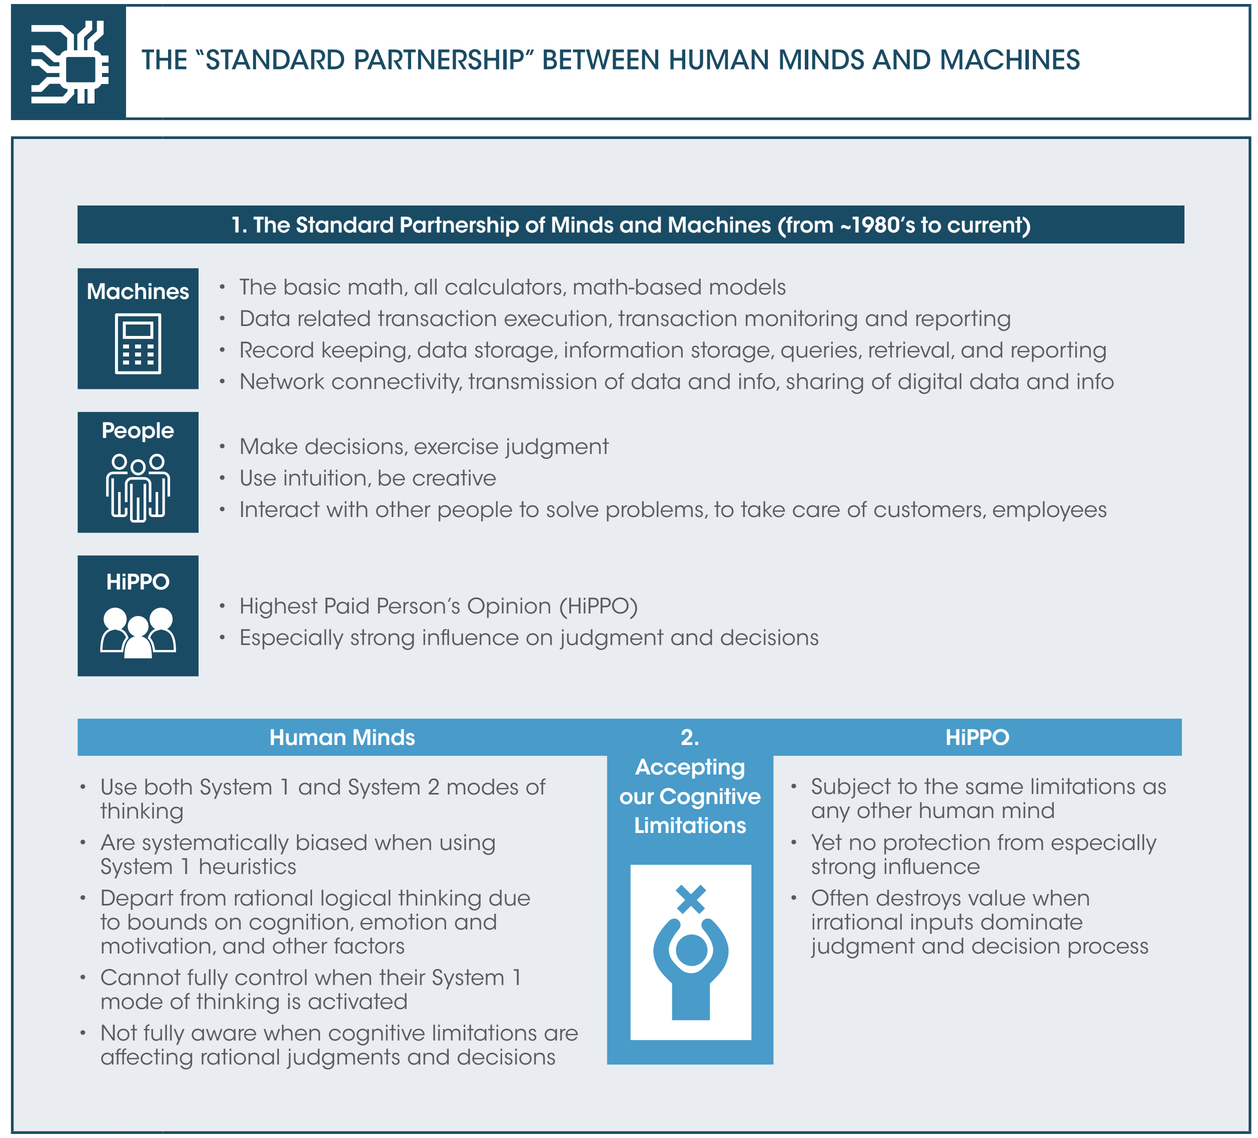

FIGURE 7 *Source: Author’s adaptation from the work of McAfee and Brynjolfsson (2017)*

model-based judgement and decision making. The psychologists Keith Stanovich and Richard West in the journal, “Behaviour and Brain Science” first introduced the specific terminology of System 1 and System 2 modes of human thinking in 2000.[21] The Stanovich and West article that first used this terminology was a part of a vast body of scientific research tracing back to the 1970s (and even earlier) that has focused on understanding the differences between so-called normative (logically rational and optimal) decision making and descriptive (actually observed) human decision making, and on explaining the reasons for why humans so often deviate from normative (and hence logically optimal) choices in their judgements and decisions.

The terminology of System 1 and System 2 modes of human thinking and the associated cognitive concepts were more widely introduced to the general public and mass media with the publication in 2011 of the book “Thinking, Fast and Slow” by Daniel Kahneman.[22] Kahneman had received the Nobel Prize in Economics in 2002 for his work with Amos Tversky (who was already deceased at that time and therefore not eligible to be a co-recipient of the Nobel Prize). The Nobel Prize committee noted that Kahneman received the Economics Prize, “...for having integrated insights from psychological research into economic science, especially concerning human judgment and decision-making under uncertainty.”[23]

I provide additional information on the descriptions of System 1 and System 2 modes of human thinking in Figures 8 and 9 below. The figures illustrate well-documented biases in human judgement and decision-making that result from the fact that in many situations, especially when there is uncertainty or when people are under time pressure or emotionally aroused, they unavoidably resort to System 1 heuristics. In their book, McAfee and Byrnjolfsson provide summaries of recent studies that demonstrate the weaknesses of System 1 thinking and intuition, even with experts.

Based on the huge body of evidence on the nature of human judgement and decision-making, and on how expert decision-making compares to the results of statistical models (when the data is available and the models can be validated), McAfee and Byrnjolfsson come to the following sobering conclusion, "...that we need to rely less on expert judgement and predictions." Their justification is as follows: "In case after case, when a model can be created and tested, it tends to perform as well as, or better than, human experts making similar decisions. Too often, we continue to rely on human judgement when machines can do better."[24]

SYSTEM 1 V. SYSTEM 2 COGNITION IN PEOPLE

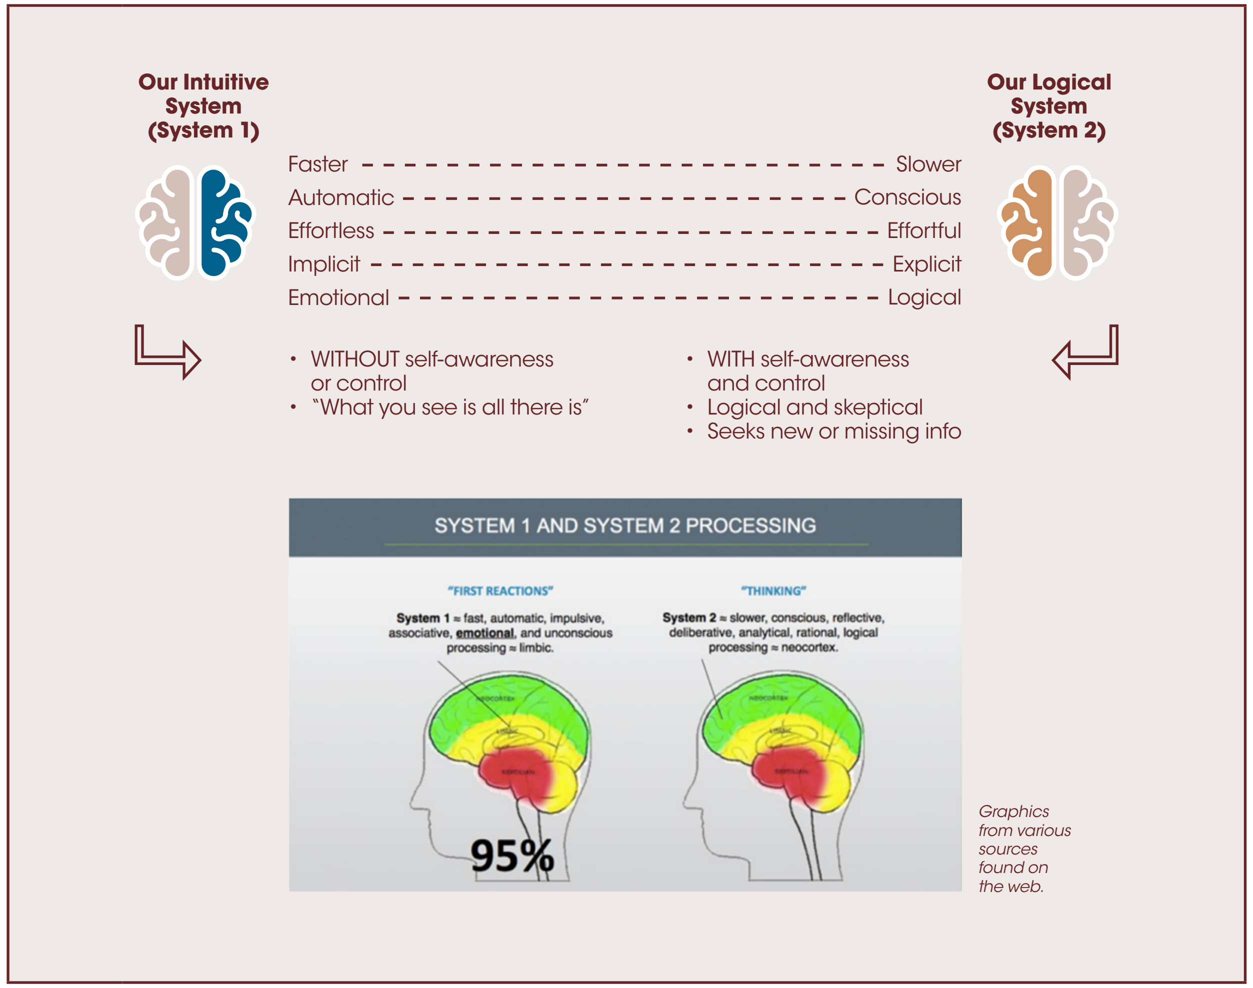

FIGURE 8

Source: Author's adaptation from the work of many researchers in the behavioural decision-making research community

INNATE PROPERTIES AND LIMITATIONS OF HUMAN JUDGMENT AND DECISION MAKING

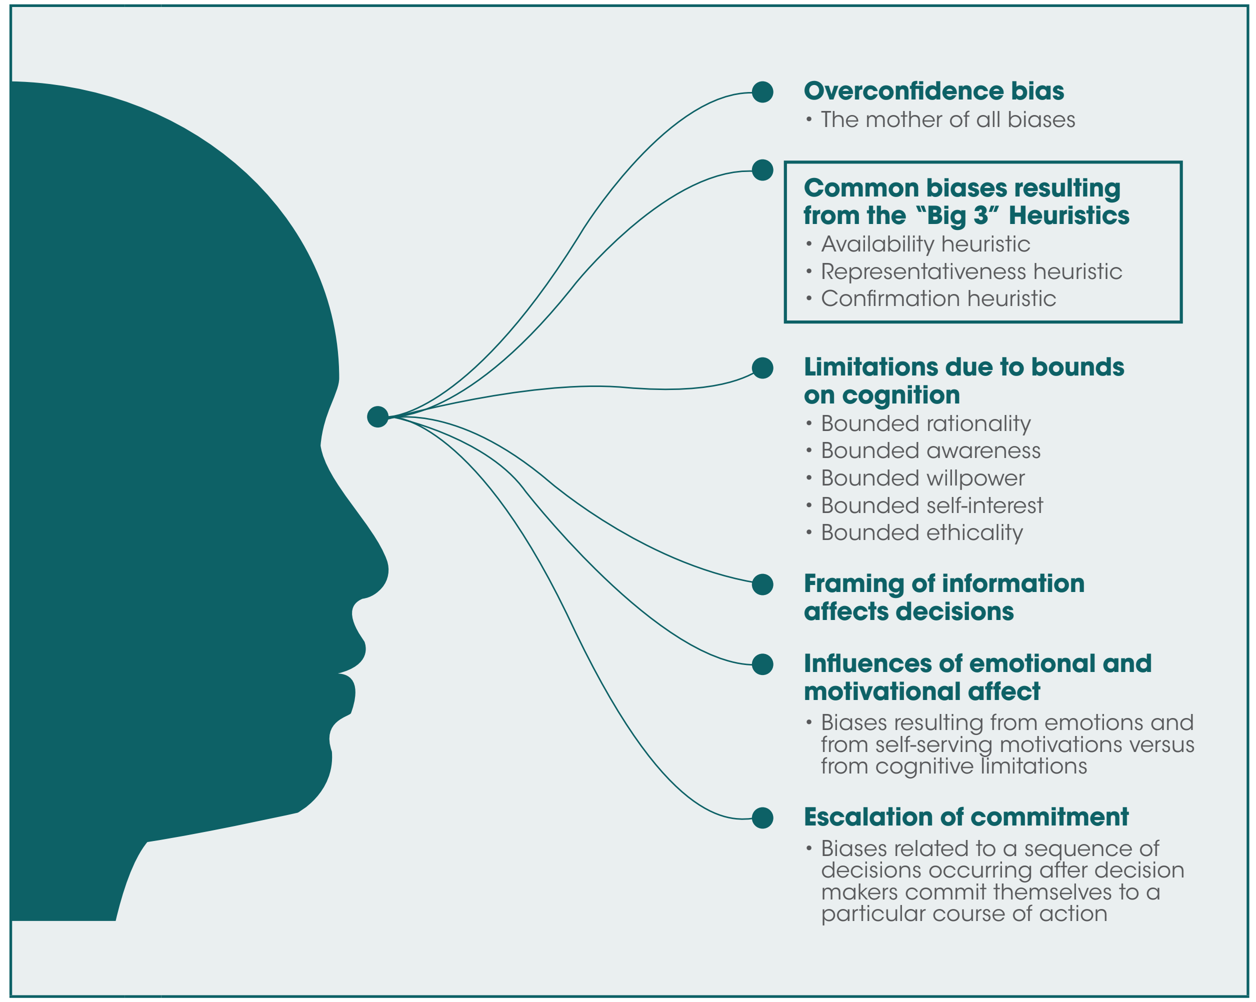

FIGURE 9

Source: Author's adaptation from the work of Bazerman and Moore (2012)[25]

McAfee and Brynjolfsson go on to raise the question about what role should people play in making decisions. It is a complicated question. Even though the results are overwhelmingly clear that statistically-based models usually outperform human expert judgement in those situations when the models and supporting data are available and when the models are properly validated, most organisations exist in dynamic, uncertain and ever-changing environments. There will always be many situations when the necessary data is not available, or where a model has not yet been created, or has not been sufficiently tested and reliably validated for a new or changed situation.

On the one hand, we know from decades of scientific research on human judgement and decision making that the System 1 mode of thinking often leads people to depart from normative logic, and often leads to questionable or outright wrong choices and decisions. On the other hand, these same cognitive mechanisms give people a range of truly remarkable capabilities. These include:

- Abilities to understand physical, social and situational context;
- Abilities to understand the intent of other people as well as the needs of other people;
- General purpose, common-sense reasoning ability that does not require domain-specific training examples;

- Associative reasoning ability based on prior (and not directly related) accumulated knowledge that can generate useful suggestions in a situation even when there is very limited data or no data at all for that specific situation;
- Abilities to identify important questions and strategic issues;
- Abilities to ask and pursue 'what if' questions and counterfactual propositions; and
- Abilities to look at situations from multiple perspectives.

With all the attention given to the impacts of technology on the business landscape, most business people forget that the great progress of building and deploying real-world AI systems in recent decades has occurred in parallel with equally great progress in understanding the nature of human judgement and decision making. McAfee and Brynjolfsson emphasise that the results that emerged from these parallel developments require organisations and society to fundamentally rethink how to put human intelligence in the loop of integrated machine and human decision making systems in an 'intelligent' way, based on the new realities of what machine intelligence is capable of, and our expanded understanding of the nature (strengths and limitations) of human intelligence.

They cite examples that lead them to endorse, "...the wisdom of having human judgement and algorithms work together."[26] McAfee and Brynjolfsson strongly support having humans and machines augment one another. They mention that some leading edge companies are already experimenting with an 'inverted partnership' between human minds and machines driven by data and algorithms. The inversion is as follows: Rather than the prior standard partnership model of having the machine provide data as an input to human judgement, the inversion is to have the human provide their judgement and intuition as an input into the machine's algorithm.

McAfee and Brynjolfsson see the broad direction for integrating human minds with increasingly intelligent machines as follows: Let algorithms and computer systems make the decisions when possible to do so (when there is the necessary data, models, and rigorous validation), sometimes with human judgement as an input. Let people override the algorithm and computer-made decisions when deemed appropriate to do so due to unusual, special or new conditions. They emphasise the importance of holding *both* the humans and machines accountable for their judgements and decisions. They emphasise it is important to keep score of the quality of decision making for the automated system and for the human working with the system.

The recent studies by Davenport, McAfee and Brynjolfsson, and Daugherty and Wilson all emphasise the importance of understanding how AI can be used for augmenting human intelligence and capabilities across all business settings. All three studies are telling us to get ready for big changes that are related to the increasing use of AI-enabled systems and machines within every type of workplace. All three studies provide evidence that these changes are already underway, and make it clear we are just at the beginning of what will be a long and profound transformation period.

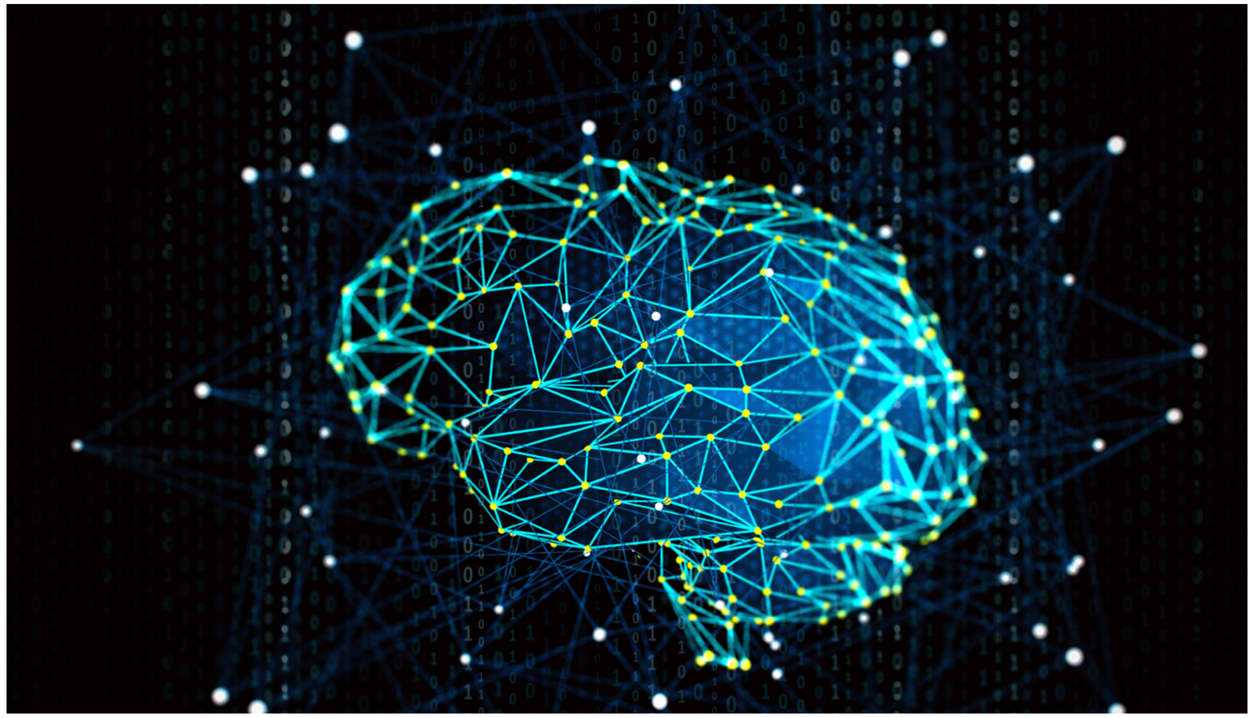

Perhaps there are advantages for the vast majority of ASEAN and Asian companies that have not yet started down this transformation pathway. They can start now, and take advantage of the lessons learned to date from the experiences of the many firms in North America and Europe, as well as from the Chinese AI superpower companies, that have already started on this transition. They can benefit from the findings of the studies summarised in this article. Also, the firms in this region who are just starting now can benefit from the fact that the AI software and solution vendor community has more experience, and more robust product and service offerings.

In the discussions above on how the partnerships and nature of symbiosis between human minds and computer-based machines needs to change, the emphasis has been on office, administrative and commercial types of work. The same point has to be extended to the many other types of work settings where decision-making occurs in the context of physical interactions and physical service delivery. This includes many different types of service work, field-work, transportation work, infrastructure related construction and maintenance work, manufacturing work, and logistics work. Similarly, the discussion above on how to put human intelligence in the loop in intelligent ways, in partnership with increasingly intelligent machines, emphasised office, administrative and commercial work. This same point has to be extended to other types of work settings where decision-making occurs in the context of physical interaction and physical service delivery.

In essence, this 'Second Machine Age', a term MacAfee and Brynjolfsson coined in their prior book[27] will propagate across all sectors of the economy and will result in most types of work being transformed in one way or another.[28]

The imperative of a new human-machine symbiosis

The earlier waves of machine automation were built around economies of scale. Automation was used to increase efficiency and raise productivity. Previous generations of automation, including computerisation, were very limited in intelligence, and as such, relatively inflexible. As a result, a requirement for automating and computerising in prior machine ages, up to very recently, has been the strict enforcement of standardisation combined with related efforts to constrain allowable change and adaptation. In other words, automation and computerisation gave us efficiency and productivity but at the cost of constraining our ability to change and adapt. Companies like Toyota and Exxon taught the post-World War II world how to use automation and computerisation to take advantage of economies of scale across the full spectrum of the product realisation life cycle, from earlier stage product development, into manufacturing with all of its required supplied chains, and through final delivery via global distribution channels. Toyota and Exxon achieved significant positions and domination in their respective global markets for decades through their ability to master and exploit economies of scale. While they respectively introduced new products periodically in response to market change, they had to carefully constrain and manage the rate and extent of change as it was so expensive and time consuming to re-engineer their highly standardised, automated and inflexible processes.

We have now entered a new machine age characterised by AI-enabled automation and computerisation, and the fusion of AI abilities with a wide range of other evolving technologies. MacAfee and Brynjolfsson refer to this new period as 'The Second Machine Age', and more commonly across industry, this new period is referred to as 'Industry 4.0', or the Fourth Industrial Revolution. This new period of automation and computerisation is built around the foundations of the economies of scalable learning versus the traditional economies of scale. Google and Amazon have become exemplars of globally prominent and dominant firms that have demonstrated the capability to put economies of scalable learning into large scale practice across the end-to-end service realisation and delivery life cycle. Google and Amazon, and the Chinese Internet and e-commerce giants Alibaba, Baidu and Tencent, and similar (if smaller) new firms are demonstrating the new realities and power of scalable learning as Toyota, Exxon and their peer firms previously demonstrated the realities and power of supply chain and manufacturing execution that were so important to the prior industrial age.

Closed loop feedback systems are essential to both economies of scale and economies of scalable learning. The most widely used management framework for closed loop feedback systems is the Plan-Do-Check-Act (PDCA) cycle that grew out of the work of statistical quality control pioneer Walter Shewhart (1891–1967). The PDCA cycle was popularised and globally promoted by W. Edward Deming (1900–1993) who became the world's most prominent advocate for the concept of quality and the management of quality improvement. Later in his life, Deming revised the name of the cycle to Plan-Do-Study-Act (PDSA) to emphasise the need for deeper study and analysis versus simpler and more superficial checks.

Advanced analytics and its extension into current AI applications can be used to improve an organisation's ability to execute closed loop feedback systems. AI-enabled applications and systems, together with supporting technological developments, have made it possible to dramatically increase the speed, the scale, and the accuracy at which an organisation executes closed loop feedback. This new wave of AI systems has improved an organisation's ability to use data to make predictions and has substantially reduced the costs of making predictions.[29] This new supercharged capability of executing the PDSA type of closed loop feedback system, combined with this ability to make better predictions at a fraction of the cost, gives senior managers a powerful opportunity to rethink their approach to organisational learning.[30] As Google, Amazon and the Chinese Internet giants have so clearly demonstrated, the company that can best use data at speed and scale to sense its external and internal environment, to make sense of the data, to identify what needs to happen, and to predict what will happen, will move faster and will win. In short, the companies that master economies of scalable learning will increasingly dominate the competitive landscape.[31] The companies that can best use analytics and AI to automate and augment their closed loop feedback systems will be the ones that most effectively master the economies of scalable learning.

How do we rethink our partnership between human minds and machine minds to accomplish this? The answer cannot be through automation that is mainly focused on displacement of human labour and minds.

Profitability requires growth *and* productivity. Growth requires change, and adaptation to uncertain and rapidly evolving environments. Productivity requires efficiency, and standardisation and stabilisation by reducing variances from process and environmental changes. Humans by nature are flexible, adaptable and dynamic, though not necessarily consistent or efficient. Machines and automated systems by design are highly efficient and capable of consistently performing at high levels of productivity, though limited in their ability to understand uncertainty and respond to rapidly changing environments.

The imperative for every organisation in this new age of AI-enabled smart machines is to find the right combination of using highly adaptable and versatile human employees with highly efficient and consistent machines that are increasingly intelligent, though in limited ways. Even with the continuation of remarkably rapid advances in AI technologies and applications, humans will still be more versatile, more adaptable, and more capable of comprehending and interpreting unquantifiable contexts and more capable of generating ideas and possibilities that do not require prior experience or available data.

Humans are key to our ongoing ability to thrive on change. Intelligent machines are key to our ability to move forward at speed, scale, and at high levels of productivity. We need to simultaneously respond to increasing rates of change in beneficial ways and also achieve high levels of efficiency and productivity. Going forward, we need a new level of human-machine symbiosis so we can leverage this new generation of intelligent machines to augment the innate and remarkable capabilities of humans in ways that enhance business and organisational capabilities for both adaptation and productivity.

Steven M. Miller
is the Vice Provost (Research) and Professor of Information Systems (Practice) at Singapore Management University. He has worked on applications and implications of advanced information systems and automation over the past four decades as an academic investigator, an industry practitioner, an industry consultant, and a university-based research planner and strategist.

References

[1] Thomas H. Davenport and Rajeev Ronaki, "Artificial Intelligence for the Real World: Don't start with Moonshots," Harvard Business Review, January-February 2018, pages 108-116. This work is substantially expanded upon and further explained in the forthcoming book, Thomas H. Davenport, "The AI Advantage: How to Put the Artificial Intelligence Revolution to Work", MIT Press, forthcoming October 2018.

[2] Ibid.

[3] Deloitte, "Deloitte State of Cognitive Survey," August 2017. Some of the results of this Deloitte survey are included in the Davenport and Ronanki's 2018 Harvard Business Review article, as Davenport was one of the lead authors on the report summarising this Deloitte survey, and Ronanki was a contributor.

[4] The Economist Intelligence Unit, "Artificial Intelligence in the Real World, "January 2017. This survey is based on interviews with 203 executives worldwide completed in 2016, of which 35 percent of respondents were from Asia. The EIU survey team constructed an AI Implementation score for Asia, Europe and North American based on their interviews. The implementation score was lowest for Asia, where implementation was in a more nascent stage compared to the other regions. They found there was a higher degree of AI-related exploration, experimentation, application and widespread deployment in North America and in Europe compared to Asia. While there are several well-known AI powerhouses in China (Alibaba, Baidu, Tencent), the EIU survey reflects the broader pattern of AI adoption across all sectors of the economy in Asia.

[5] For example, in Singapore, SMEs, defined as those with annual revenue of S$100M or less, account for 60 percent of total national employment, and just over one-third of gross domestic product. According to the official ASEAN SME website (www.aseansme.org), SMEs in ASEAN member states account for between 51 to 97 percent of national employment, and between 23 to 58 percent of gross domestic product.

[6] Thomas H. Davenport and Julia Kirby, "Only Humans Need Apply: Winners & Losers in the Age of Smart Machines," Harper Business, 2016.

[7] All illustrations attributed to Steven Miller (adapted from the published content of other authors) were created with the substantial assistance and contributions of Justin Lebrun of Scientific Reach, http://www.scientificreach.com/.

[8] "Insurance claims assessment in seconds soon a reality for Prudential Singapore PRUshield policyholders," Prudential website, November 13, 2017.

[9] Thomas H. Davenport and Julia Kirby, "Beyond Automation," Harvard Business Review, June 2015.

[10] J. C. R. Licklider, "Man-Computer Symbiosis," IRE Transactions on Human Factors in Electronics, March 1960, volume HFE-1, pages 4-11.

[11] Douglas C. Engelbart, "Augmenting Human Intellect: A Conceptual Framework," October, 1962, Stanford Research Institute, Summary Report AFOSR-3223. This report was prepared for the Director of Information Sciences, U. S. Air Force Office of Scientific Research, Washington DC.

[12] When Englebart submitted his report-cum-proposal to the U.S. Department of Defense, J. C. R. Liklider had recently assumed a position at Department of Defense's Advanced Research Projects Agency (ARPA), with responsibility for giving out grant money to support developments in the computing area. Liklider decided that his unit should fund Englebart's proposal. Examples of innovations that grew out of Englebart's R&D programme on Augmented Intelligence are described at http://dougengelbart.org/library/engelbart-archives.html#Firsts.

[13] An excellent summary of parallel trajectories of development of AI for replicating and replacing the human versus AI for augmenting the human are given in the book by Walter Issacson, "The Innovators: How a group of hackers, geniuses and geeks created the digital revolution," Simon & Schuster, 2014, see Ch 12, Ada Forever.

[14] Paul Daugherty and Jim Wilson, "Humans + Machine: Reimagining Work in the Age of AI," Harvard Business Review Press, March 2018.

[15] According to Daugherty and Wilson, the first wave of business process transformation occurred around 1908 with the Ford Motor Company's production of the Model T using standardised assembly lines. This was the beginning of the age of standardised mass production. The second wave started in the 1970's, continuing into the 1980's and 1990's with the widespread computerisation of work processes that occurred at that time as a result of the diversification of computers (mini computers, PCs, more distributed computing) beyond mainframes, and the emergence of various types of enterprise software systems. They refer to the current and emerging business process transformation efforts as the third wave, with the key distinction of this newest wave being the ability to create processes that are flexible, adaptable and not always predetermined, while at the same time being scalable, efficient and repeatable.

[16] Paul Daugherty and Jim Wilson, "Humans + Machine: Reimagining Work in the Age of AI," Harvard Business Review Press, March 2018.

[17] Ibid.

[18] Ibid.

[19] The two other books Andrew McAfee and Erik Brynjolfsson published together were, "The Race Against the Machine," Digital Frontier Press, 2012, and "The Second Machine Age," W. W. Norton, 2016.

[20] Andrew McAfee and Erik Brynjolfsson, "Machine Platform Crowd: Harnessing our Digital Future," W.W. Norton and Company, 2017.

[21] Keith E. Stanovich and Richard F. West, "Individual differences in reasoning: Implications for the rationality debate," Behaviour and Brain Sciences, Volume 23, year 2000, pages 645–726.

[22] Daniel Kahneman, "Thinking, Fast and Slow", Farrar, Straus and Giroux, 2011.

[23] "Daniel Kahneman: Facts", Nobelprize,org, https://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/2002/kahneman-facts.html .

[24] Andrew McAfee and Erik Brynjolfsson, "Machine Platform Crowd: Harnessing our Digital Future," W.W. Norton and Company, 2017.

[25] Max H. Bazerman and Dan A. Moore, "Judgement in Managerial Decision Making, 8th Edition", John Wiley & Sons, 2012.

[26] Andrew McAfee and Erik Brynjolfsson, "Machine Platform Crowd: Harnessing our Digital Future," W.W. Norton and Company, 2017.

[27] Andrew McAfee and Erik Brynjolfsson, "The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies," W.W. Norton and Company, 2016.

[28] An excellent overview on the impacts of IT, AI and Automation on the workforce is provided in U.S. National Academies of Sciences, Engineering, and Medicine, "Information Technology and the U.S. Workforce: Where Are We and Where Do We Go from Here?" The National Academies Press, 2017, doi:10.17226/24649.

[29] Ajay Agrawal, Joshu Gans and Avi Goldfarb, "Prediction Machines: The simple economics of artificial intelligence," Harvard Business School Press, 2018.

[30] See Linda Argote, "Organizational Learning: Creating, Retaining and Transferring Knowledge, 2nd Edition", Springer Science + Business Media, 2013, for a comprehensive overview of academic research on learning at the organisational level. In this book, Prof Argote examines evidence that shows that organisations vary tremendously in the rate at which they learn, and she reviews and explains the academic research available as of that point in time on factors that explain the variance observed in organisational learning curves.

[31] See John Seely Brown's recent discussions on scalable learning. For example, John Seely Brown, "Working/Learning/Leading in the Exponential Age", Adapted from the plenary session given at the 2016 AACSB International Deans Conference, July 2016; and John Seely Brown, "Sense-making in our post AlphaGo World", Stanford MediaX Keynote 2017, April 20, 2017.