12-2021

AI and the future of work: What we know today

Steven M. MILLER
Singapore Management University, stevenmiller@smu.edu.sg

Thomas H. DAVENPORT
Babson College

AI and the Future of Work: What We Know Today

Steven M. Miller and Tom Davenport

Published in The Gradient, 18 December 2021

https://thegradient.pub/artificial-intelligence-and-work-two-perspectives/

Table of Contents[1]

Part I: Impacts on Jobs and the Nature of Work

One of the most important issues in contemporary societies is the impact of automation and intelligent technologies on human work. Concerns with the impact of mechanization on jobs and unemployment go back centuries, at least since the late 1500's, when Queen Elizabeth I turned down William Lee's patent applications for an automated knitting machine for stockings because of fears that it might turn human knitters into paupers.[2] In 1936, an automotive industry manager at General Motors named D.L. Harder coined the term "automation" to refer to the automatic operation of machines in a factory setting. Ten years later, when he was a Vice President at Ford Motor company, he established an "Automation Department" which led to widespread usage of the term.[3]

The origins of intelligent automation trace back to US and British advances in fire-control radar for operating anti-aircraft guns to defend against German V-1 rockets and aircraft during World War II. After the war, these advances motivated the MIT mathematician Norbert Weiner to develop the concept of "cybernetics", a theory of machines and their potential based on feedback loops, self-stabilizing systems, and the ability to autonomously lean and adapt behavior.[4] In parallel, the Dartmouth Summer Research Project on Artificial Intelligence workshop was held in 1956 and is recognized as the founding event of artificial intelligence as a research field.[5]

Since that decade, workplace automation, cybernetic-inspired advanced feedback systems for both analogue and digital machines, and digital computing based artificial intelligence (together with the overall field of computer science) have advanced in parallel and co-mingled with one another. Additionally, opposing views of these developments have co-existed with one side highlighting the positive potential for more capable and intelligent machines to serve, benefit and elevate humanity, and the other side highlighting the negative possibilities and threats including mass unemployment, physical harm and loss of control. There has been a steady stream of studies from the 1950's to the present assessing the impacts of machine automation on the nature of work, jobs and employment, with each more recent study considering the capability enhancements of the newest generation of automated machines.

To contribute to a better understanding of the contemporary realities of AI workplace deployments, the two of us (Davenport and Miller) recently completed 29 case studies of people doing their everyday work with AI-enabled smart machines.[6] Twenty-three of these examples were from North America, mostly in the US. Six were from Southeast Asia, mostly in Singapore. In this essay, we compare our findings on job and workplace impacts to those reported in the MIT Task Force on the Work of the Future report, as we consider that to be the most comprehensive recent study on this topic.

MIT established its Work of the Future Task Force in 2018 as an "institute-wide initiative to understand how emerging technologies are changing the nature of human work and the skills required—and how we can design and leverage technological innovations for the benefit of everyone in society."[7] The task force focused on understanding the current and forthcoming impacts of advanced automation—in particular, artificial intelligence and robotics—on the nature of work, on productivity and jobs, and on labor markets and employment trends. Their final report was published in November 2020 and mostly focused on the situation in the US, though their field studies also included visits to German factories. They also extensively reviewed research studies on the workforce, employment, and labor market impacts of automation—with emphasis on impacts of AI and robotics—from all over the world.[8] The task force effort included in-depth field studies in five industry areas: insurance, healthcare, vehicle driving (autonomous vehicles), warehousing and logistics, and manufacturing.[9]

Our case studies also included examples from insurance, healthcare, and manufacturing settings, as well as from various other service sector settings, other production operation settings, and field work settings for public safety and infrastructure operations. A listing of our case studies organized by the functional areas that the AI system is supporting is shown in Table 1 below.

We compare three of the six major conclusions extracted from the MIT task force final report with our case study findings where our study efforts overlap. In the first two areas, the task force's conclusions are entirely consistent with what we found. In the third area we observed some differences between the MIT study's findings and our own. We conclude with brief comments on the three other MIT Task Force conclusions that were beyond the scope of our study effort because we feel that these other national level policy issues are important for readers of this essay to be aware of. Quotations colored in blue are directly extracted from the MIT Work of the Future task force reports.

AI system is used to support the following functional area in each case	Case Name
Sales and Business Development	Morgan Stanley: Financial Advisors and The Next Best Action System
	ChowNow: Growth Operations and RingDNA
	Stitch Fix: AI-Assisted Clothing Stylists
	Arkansas State University: Fundraising with Gravyty
Product Development Management	Shopee: The Product Manager's Role in AI Driven E-Commerce
Administrative Operations	Haven Life & Mass Mutual: The Digital Life Underwriter
	Radius Financial Group: Intelligent Mortgage Processing
	DBS Bank: AI-Driven Transaction Surveillance
	Medical Diagnosis and Treatment Record Coding with AI
	Dentsu: RPA for Citizen Automation Developers
IT and analytics support	84.51° & Kroger: AutoML To Improve Data Science Productivity
	Mandiant: AI Support for Cyber Threat Attribution
Customer and product support	DBS Digibank India: Customer Science for Customer Service
	Intuit: AI-Assisted Writing with Writer.Com
	Lilt: The Computer-Assisted Translator
Governance and Ethics	Salesforce: Architects of Ethical AI Practices
Professional Services (medical, legal)	The Dermatologist: AI-Assisted Skin Imaging
	Good Doctor Technology: Intelligent Telemedicine in Southeast Asia
	Osler Works: The Transformation of Legal Service Delivery
Manufacturing and Other Production Operations	PCB Linear: AI Enabled Virtual Reality for Employee Training
	Seagate: Improving Automated Visual Inspection of Wafers and Fab Tooling with AI
	Stanford Health Care: Robotic Pharmacy Operations
	Fast Food Hamburger Outlets: Flippy--Robotic Assistants for Fast Food Preparation
	FarmWise: Digital Weeders for Robotic Weeding of Farm Fields
Public Safety and Infrastructure Operations	Wilmington, North Carolina Police Department: AI Driven Policing
	Certis: AI Support for the Multi-Faceted Security Guard at Jewel Changi Airport
	Southern California Edison: Machine Learning Safety Data Analytics for Front Line Accident Prevention
	Mass Bay Transit Authority: AI Assisted Diesel Oil Analysis for Train Maintenance
	Singapore Land Transport Authority: Rail Network Management in a Smart City

Table 1. Case studies in the forthcoming Davenport/Miller book, "Working with AI: Real Stories of Human Machine Collaboration", forthcoming, MIT Press 2022.

Technology Is Not Replacing Human Labor En Masse Anytime Soon

The first MIT task force conclusion addresses whether technology will replace human labor:

Technological change is simultaneously replacing existing work and creating new work. It is not eliminating work altogether.

No compelling historical or contemporary evidence suggests that technological advances are driving us toward a jobless future. On the contrary, we anticipate that in the next two decades, industrialized countries will have more job openings than workers to fill them, and that robotics and automation will play an increasingly crucial role in closing these gaps.

Their report acknowledges that intelligent machines are thus far capable of completing particular tasks. In most cases they cannot perform entire jobs, and are seldom able to automatically perform entire business processes. This makes it very unlikely that large-scale automation of human labor will take place over the next few decades. Indeed, in all of our case studies, the organizations involved said that AI and robotics had freed up workers to perform more complex tasks, and human workers had not lost jobs because of automation or AI. Several of the jobs we described across our collection of case examples are new and wouldn't exist without AI. Many of the companies we profiled were growing (in part because of their effective use of digital and AI technologies), so they needed all their human workers to keep up with growth.

The MIT task force report highlights that from a broad economic perspective, growth of economies, demographics, and restrictive immigration policies will make it is far more likely that many jobs will go unfilled over the next few decades because labor is in short supply, at least in most of the world's largest economies. World Bank statistics point in the same direction as this assessment, indicating that in 11 of the world's 12 largest economies, fertility rates (births per woman) have been well below replacement levels and the proportion of the population age 65 and over has been on an increasing trajectory.[10] The inevitable implication is that at least in these 11 economies that are currently the world's largest, human labor will increasingly be in short supply.

If industrialized countries will have more job openings than workers to fill them even with increasing workplace usage of AI and robotics and other types of technologies- as predicted by the MIT report- this suggests that signs of this trend should already be visible in countries where labor is already in especially short supply. In fact, recent work by the economists Daron Acemoglu and Pascual Restrepo provides evidence that, "Indeed, automation technologies have made much greater inroads in countries with more rapidly aging populations," and that "the adoption and development of these technologies are receiving a powerful boost from demographic changes throughout the world and especially from rapidly-aging countries such as Germany, Japan and South Korea."[11]

These reasons explain why the MIT task force report forecasts that neither the US nor the world at large is heading towards a future where there is not enough work for people to do as a result of greater usage of more sophisticated automation. More likely, in the decades to come, most of the world's largest economies will make even greater use of AI, robotics and other existing types of automation to keep their economic output from shrinking given their slowing or even declining labor force participation rates. The remaining human labor will be indispensable in making this transition, though of course people will need the education and upskilling required to participate in this effort.

Organizational Changes from AI Are Happening Gradually

The second task force conclusion sheds light on the confusing dichotomy between the rapid pace of AI technology development as viewed from R&D and tech start-up announcements and the much slower pace at which organizations are able to absorb and productively harness AI and robotic capabilities. It is described here:

Momentous impacts of technological change are unfolding gradually.

Spectacular advances in computing and communications, robotics, AI, and manufacturing processes are reshaping industries as diverse as insurance, retail, healthcare, manufacturing, and logistics and transportation. But we observe substantial time lags, often on the scale of decades, from the birth of an invention to its broad commercialization, assimilation into business processes, widespread adoption, and impacts on the workforce … Indeed, the most profound labor market effects of new technology that we found were less due to robotics and AI than to the continuing diffusion of decades-old (though much-improved) technologies of the Internet, mobile and cloud computing, and mobile phones. This timescale of change provides the opportunity to craft policies, develop skills, and foment investments to constructively shape the trajectory of change toward the greatest social and economic benefit.

Across our 29 case studies, we also observed that new AI-based systems, their supporting platforms and infrastructure, and their surrounding work processes do not materialize easily or quickly. It takes time for an organization to orchestrate the deep collaborations and complex deployment efforts across the ecosystem of job roles that need to be involved within the company internally, and also within key external partner organizations (vendors, and sometimes customers).[12] Most of our case examples were the result of multi-year AI deployment and related process improvement efforts that started a well before we interviewed system users and other company personnel. Our interviews occurred after the companies had started to realize tangible improvements in efficiency and effectiveness after deploying and stabilizing a new AI system.

Major process changes, especially those involving AI systems, require up-front and ongoing investments, not only in the direct software and hardware aspects of the technology, but also in complementary enabling efforts (e.g., data acquisition, data engineering, infrastructure enhancement, new types of testing and validation efforts) and organizational adjustments (e.g., policy and process changes, new types of reviews for bias, fairness and other aspects of responsible AI usage) required to harness the new capabilities. On top of this, we learned from some of our case studies that managing the organizational change effort, educating the relevant parts of the workforce about using a new AI model, and gaining employee trust to use the new AI-enabled systems in their everyday work can sometimes take far longer than developing and testing the model.

Indeed, new AI developments are proceeding at breakneck speed, but bringing everything together across technology, people, and job roles in any real-world work setting is a very complex, time intensive and iterative undertaking that extends over longer time periods.

The MIT task force elaborated on this slow adaptation process:

As this report documents, the labor market impacts of technologies like AI and robotics are taking years to unfold … in each instance where the Task Force focused its expertise on specific technologies, we found technological change — while visible and auguring vast potential — moving less rapidly, and displacing fewer jobs, than portrayed in popular accounts. New technologies themselves are *often astounding, but it can take decades from the birth of an invention to its commercialization, assimilation into business processes, standardization, widespread adoption, and broader impacts on the workforce*.[13]

The "Productivity J-Curve" phenomenon described by Professor Erik Brynjolfsson and his colleagues[14] provides a conceptual framework and explanation for why the observed rate of AI and robotics assimilation within a specific company is a slow and gradual process. In their research brief prepared for the MIT task force, they described the productivity J-curve phenomenon as follows:

… new technologies take time to diffuse, to be implemented, and to reach their full economic potential. For a transformative new technology like AI, it is not enough to simply "pave the cow paths" by making existing systems better. Instead, productivity growth from new technologies depends on the invention and implementation of myriad complementary investments and adjustments. The result can be a productivity J-curve, where productivity initially falls, but then recovers as the gains from these intangible investments are harvested.

Productivity growth is the most important single driver of higher living standards, and technological progress is the primary engine of productivity growth. Thus, it is troubling that despite impressive advances in AI and digital technologies, measured productivity growth has slowed since 2005.

While there are many reasons for this, the most important is that technological advances typically don't translate into improvements in productivity unless and until complementary innovations are developed. These include many intangible assets such as new business processes, business models, skills, techniques, and organizational cultures. The need for myriad complementary innovations is substantial, especially in the case of fundamental technology advancements such as AI. Yet, these complementary innovations can take years or even decades to create and implement; in the meantime, measured productivity growth can fall below trends as real resources are devoted to investments in these innovations. Eventually, productivity growth not only returns to normal but even exceeds its previous rates. This pattern is called a Productivity J-Curve.

Of course, a company can use a cloud-based or other AI application that does not require deep levels of integration with its existing technical infrastructure or processes. In such cases, the time span required to realize benefits could be short, and there may not be much or any productivity J-curve effect. We observed this type of situation in only two of our cases. One was with a private practice dermatologist who made use of a cloud-based AI-enabled system that his patients would also use at home so he could track the progress of high-risk dermatology cases.[15] In the second example, a company would simply send their documents requiring professional caliber translation to a cloud-based "translation-as-a-service" provider that combined human translators with AI support systems to achieve highly demanding or non-standard language translation in a way that is highly productive as well as nuanced, context-specific, and appropriately edited.[16] Such situations exist, but have an inherently smaller degree of impact on the company's productive capabilities exactly because there is no deep integration with or improvements to existing infrastructure and processes. All of our other case examples required deep integration and/or major supporting changes to their internal processes that extended over multi-year time periods.

For example, while we were preparing our case study on AI-enabled financial transaction surveillance at DBS Bank, the company's Chief Analytics Officer Sameer Gupta shared with us:

In my view, the reason this effort has been so successful is that it was not just been about analytics and AI. The team looked at how they run the entire function of transaction surveillance, transforming how they do this function end-to-end. This transformation has been supported, supplemented and augmented by analytics. But even with the best analytics models, had we not done all the other changes involved in this transformation, we would not have obtained the very impressive results that we ended up achieving. I see this as a successful business transformation that was augmented by analytics.

Sameer Gupta's comment illustrates how AI system deployments require supporting implementation of many other business and organizational adjustments. In two of our case studies, large firms purchased a subsidiary to speed up their journey of capability development: MassMutual's purchase of Haven Life for digital underwriting and Kroger's purchase of 84.51o for data science capabilities. Despite acquiring entire organizational units with strong capabilities for creating and using the AI-based systems, the two large parent firms still had to go through a multi-year process to integrate both the technical capabilities as well as the "way of working" capabilities of these newly acquired subsidiaries into their overall ecosystems.

There is no escaping the reality that it takes substantial effort over an extended period of time for a company to make the necessary supporting complimentary investments and adjustments—above and beyond the direct investments and efforts required to design, build and deploy new AI systems—to assimilate these new technologies in ways that lead to substantial increases in productivity. Senior management in both the private and public sectors overseeing investments in AI and other advanced automation projects need to understand and anticipate the extended time periods required for an organization to make the necessary complementary investments, innovations and adjustments to go beyond merely deploying the technology. They also need to anticipate the productivity J-curve effect.

But it can be worth the effort. All of our case examples provide examples of productive capacity improvements either in terms of task or process output capacity, quality, or a combination of both.

Augmentation Much More So Than Automation

The MIT report emphasizes that augmentation is both a more desirable and more common outcome than large-scale automation. Augmentation is where employers create workplaces that combine smart machines with humans in close partnerships—symbiotically taking advantage of both human intelligence and machine intelligence. In other words, the AI system is used to complement the capabilities of a human worker (or vice versa). Economists use the term "automation" to refer to situations where the deployment of a machine in the workplace (including AI systems and robots) results in direct substitution, and the human worker who was previously doing that job is replaced by the machine (and the company may or may not redeploy that worker elsewhere within the company).[17] Most of our 29 case studies were examples of augmentation, and from what we observed, AI augmentation is largely quite successful. For the few cases that involved full automation of a limited set of tasks, there was still a need for humans to supervise and support the continuous improvement of the fully automated task or process, and to handle special cases and disruptions. The fact that both our case studies and the MIT task force's field studies observed far more instances of worker augmentation than full automation is consistent with the key points above that technology is not replacing human labor en masse anytime soon, and that changes in organizational "ways of working" are happening gradually.

The MIT task force effort included an imaginative and increasingly plausible view of how augmentation can be taken to even higher levels and expand into new types of applications. These ideas come from the task force research brief on "Artificial Intelligence and the Future of Work".[18] The brief's authors emphasize "thinking less about people OR computers and more about people AND computers." They elaborated as follows:

By focusing on human-computer groups—superminds—we can move away from thinking of AI as a tool for replacing humans by automating tasks, to thinking of AI as a tool for augmenting humans by collaborating with them more effectively. As we've just seen, AI systems are better than humans at some tasks such as crunching numbers, finding patterns, and remembering information. Humans are better than AI systems at tasks that require general intelligence—including non-routine reasoning and defining abstractions—and interpersonal and physical skills that machines haven't yet mastered. By working together, AI systems and humans can augment and complement each other's skills.

The possibilities here go far beyond what most people usually think of when they hear a phrase like "putting humans in the loop." Instead of AI technologies just being tools to augment individual humans, we believe that many of their most important uses will occur in the context of groups of humans. As the Internet has already demonstrated, another very important use of information technology—in addition to AI—will be providing hyperconnectivity: connecting people to other people, and often to computers, at much larger scales and in rich new ways that were never possible before.

That's why we need to move from thinking about putting humans in the loop to putting computers in the group.

Using technology to attain new levels of collective coordination and intelligence is not at all far-fetched. We already see this occurring to some extent in real-world situations in two of our case studies. In our Singapore Land Transport Authority (LTA) Smart City rail network management example, the FASTER system predicts impending operational disturbances and supports operations center personnel in their efforts to bridge rail operators, LTA, and all relevant government authorities who would be involved in responding to any type of disruptive incident in the rail network system.[19] In our Certis Jewel Changi Airport example, an AI-enabled multi-service orchestration platform handles the consolidation and integration of all incoming video and sensor inputs and front-line worker reports from the 10 story Jewel mega-mall, generates alerts, and augments the ability of a centralized smart operations center team of humans to do complex situation assessment, response planning, and 'man-in-the-middle' coordination and communication across multiple stakeholders. These stakeholders include ground staff at Jewel mall, senior management at Certis and Jewel, and external parties including the ambulance teams, medical facilities, and the government authorities.[20] Both of these examples are in Singapore—a city-state economy and society making the future happen now. Over time, we expect to see more examples where smart-machine augmentation happens at the level of teams, departments, and entire business groups and organizations, and not just at the level of individual employees

Part II: Skills, Labor Markets, and Policy Issues

Education, Training and Skill Development

The MIT task force examined education, training, and skill development issues to meet the needs of increasing usage of AI, robotics and other technologies, as well as to address opportunity creation for middle and low wage workers. They explained that their policy focus was on education and training for adults, particularly those whose work is more vulnerable to automation. This includes those in lower-wage jobs, those whose education pathways do not include four-year college degrees, and those who are displaced mid-career. Their key conclusion pertaining to education and training is as follows:

Fostering opportunity and economic mobility necessitates cultivating and refreshing worker skills.

Enabling workers to remain productive in a continuously evolving workplace requires empowering them with excellent skills programs at all stages of life: in primary and secondary schools, in vocational and college programs, and in ongoing adult training programs.

Of course, this conclusion is applicable to all segments of the workforce of any country's economy, though it is especially important for those in the segments more vulnerable to being displaced by automation. Across the work settings of our 29 case studies, we only interviewed people who were gainfully employed, highly engaged with the technology and process changes that had successfully taken place in their work setting, and for the most part enthusiastic about working with or managing the new AI-enabled systems in their workplace. Obviously, the types of things we learned about training and skill development given this segment of people we encountered—all in relatively positive and promising employment situations-- will be quite different from recommendations focused on people in highly vulnerable segments or segments lacking promise and opportunity.

We also found that frontline workers, in order to collaborate effectively with smart machines in their work, needed new skills. However, in contrast to the MIT report, we did not find that those skills had been acquired through "excellent skills programs" sponsored by schools, colleges, and employers. Instead, most of the new skills were acquired on the job, or by employees who were personally motivated to acquire new skills on their own.

Leading higher education institutions have already started to adopt new AI-related skills programs, but many institutions still have not done so. While some progressive employers have internally implemented AI-related skills programs, many have not. As such, the majority of existing employees in most countries are largely on their own to develop these skills. Singapore is an exception due to the SkillsFuture national initiative to provide continuing education for the existing workforce, and also due to the AI Singapore educational outreach programs.

As the education and training policy recommendations of the MIT task force final report focused more on basic skill development needs for those at risk at being excluded from workforce participation or promising employment opportunities, they did not comment on the more nuanced issue of the importance of hybridized business and IT skills that we found in many of our case studies. In our examples, organizations had to deepen their internal capabilities in IT and expand into related areas for digital transformation and data science/AI. Frontline system users had to learn how to work with the systems. Supervisors and frontline managers had to work through the process changes and learn how to manage in the new setting. Technology staff had to hybridize their skills to include business and domain understanding. Business users had to develop digital thinking and technological savviness. In addition, workers needed to move into new types of roles which spanned and integrated business and technology (for example, product management, data governance, ethical AI practices).

While both self-motivated learning and IT/business hybridization are not easy to accomplish, they are relatively straightforward to do successfully for those in the workforce with the highest levels of education (undergraduate degrees and post-graduate degrees). In fact, the MIT task force report shows that in recent decades, at least in US labor markets, those in the workforce with highest levels of education have mostly done well.[21]

A Warning About Polarization of Labor Markets

Our focused set of case studies did not address long-term economic, employment and labor market issues. But the MIT Work of the Future task force analyzed US economy and labor market trends over decades leading up to the present, highlighting the stark realities of employment polarization and diverging job quality. They spotlighted the decline in the proportion of "middle-skill jobs" in the US labor market and the fact that wages for those in low-skilled occupations have stagnated for several decades. The task force explained the situation as follows[22]:

This ongoing process of machine substitution for routine human labor tends to increase the productivity of educated workers whose jobs rely on information, calculation, problem-solving, and communication — workers in medicine, marketing, design, and research, for example. It simultaneously displaces the middle-skill workers who in many cases provided these information-gathering, organizational, and calculation tasks. These include sales workers, office workers, administrative support workers, and assembly line production positions.

Ironically, digitalization has had the smallest impact on the tasks of workers in low-paid manual and service jobs, such as food service workers, cleaners, janitors, landscapers, security guards, home health aides, vehicle drivers, and numerous entertainment and recreation workers. Performing these jobs demands physical dexterity, visual recognition, face-to-face communications, and situational adaptability, which remain largely out of reach of current hardware and software but are readily accomplished by adults with modest levels of education. As middle-skill occupations have declined, manual and service occupations have become an increasingly central job category for those with high school or lower education. This polarization likely will not come to a halt any time soon.

The task force's observation that US labor market employment polarization has been the status quo situation for over four decades now—and that the degree of polarization is more extreme in the US than in other advanced economies that have experienced positive productivity growth over past decades—led to their three additional conclusions:

Rising labor productivity has not translated into broad increases in incomes because labor market institutions and policies have fallen into disrepair.

Improving the quality of jobs requires innovation in labor market institutions.

Investing in innovation will drive new job creation, speed growth, and meet rising competitive challenges.

We feel these additional national-level policy conclusions made by the task force are important to highlight here for readers of this essay. These additional three conclusions, when combined with the other conclusions discussed above, set the stage for what is perhaps the strongest statement in their final report[23]:

Yet, if our research did not confirm the dystopian vision of robots ushering workers off of factory floors or artificial intelligence rendering superfluous human expertise and judgment, it did uncover something equally pernicious: Amidst a technological ecosystem delivering rising productivity, and an economy generating plenty of jobs (at least until the COVID-19 crisis), we found a labor market in which the fruits are so unequally distributed, so skewed toward the top, that the majority of workers have tasted only a tiny morsel of a vast harvest.[24]

These conclusions are the foundations of important warning statements made by the MIT task force team that need to be heeded by senior managers, C-suite executives and board of director members in the private sector as well as by civil servants and elected government officials. Even though their statements are directly aimed at the situation in the US, the threats associated with excluding major segments of the workforce from sharing the fruits of productivity improvement and wealth creation apply to managers and government officials in all countries. The task force's final report cautioned[25]:

Where innovation fails to drive opportunity, however, it generates a palpable fear of the future: the suspicion that technological progress will make the country wealthier while threatening livelihoods of many. This fear exacts a high price: political and regional divisions, distrust of institutions, and mistrust of innovation itself.

The last four decades of economic history give credence to that fear. The central challenge ahead, indeed the work of the future, is to advance labor market opportunity to meet, complement, and shape technological innovations. This drive will require innovating in our labor market institutions by modernizing the laws, policies, norms, organizations, and enterprises that set the "rules of the game."

Part III: Conclusion

The value of our 29 case studies summarized in Table 1 is that they provide real-world examples in every-day operational work settings of how people are already successfully collaborating with smart machines to improve business capabilities. This is not the future of work. It is already happening now in a growing number of organizations.

There is no doubt that AI is becoming both more pervasive and more capable and able to support more types of tasks. The universe of industries and jobs that already make use of AI as part of daily work is large and is growing rapidly. We foresee that in the coming years, many more workers will be asked or even required to work with smart machines. We suspect doing so would enhance their employability while refusing to do so would hinder their employment prospects.

Our examples, along with the conclusions and supporting field studies of the MIT Work of the Future Task Force report, are also a counter to the doom-and-gloom view that AI will destroy jobs. It is definitely changing work, but not destroying it.

As AI and other forms of advanced automation continue to diffuse across an entire economy, there are other aspects of the story that go beyond our case study documentation of successful and positive company efforts to combine human capabilities and machine capabilities to improve business performance in their workplace. The MIT Work of the Future task effort provides a broader view of these changes by illuminating the multiple sides of this unfolding journey from an economy-wide employment and labor market perspective. They conclude that we must drive forward with innovation- including the increased usage of AI and robotics- in order to create the new products, services and industries which lead to new job opportunities for all segments of the workforce, not just for those at the highest levels of attainment for income and education. Their conclusions also highlight the risks and perils of failing to advance labor market opportunity in light of the persistent labor market polarization, especially in the US. Rectangle

References

[1] Our original essay on "AI and Jobs: Two Perspectives" was published on the AI Singapore website on 02 September 2021, https://aisingapore.org/2021/09/artificial-intelligence-and-work-two-perspectives/. This is a substantially expanded and revised version prepared for The Gradient with guidance from the Gradient editorial staff.

[2] See Encyclopedia Britannica, "William Lee, English Inventor," https://www.britannica.com/biography/William-Lee and the Wikipedia entry on William Lee, https://en.wikipedia.org/wiki/William_Lee_(inventor).

[3] Katsundo Hitomi, "Automation-its concept and a short history," Technovation, 14(2), 1994.

[4] Thomas Rid, Rise of the Machines: A Cybernetic History, New York, W.W. Norton & Company, 2016.

[5] See the Wikipedia entry on the Dartmouth workshop, https://en.wikipedia.org/wiki/Dartmouth_workshop

[6] Thomas H. Davenport and Steven M. Miller. Working with AI: Real Stories of Human-Machine Collaboration. MIT Press, forthcoming in 2022.

[7] This description of the purpose of the MIT Future of Work Task Force is stated on their website homepage at https://workofthefuture.mit.edu/.

[8] David Autor, David Mindell, and Elisabeth Reynolds, "The Work of the Future: Building Better Jobs in an Age of Intelligent Machines," report published by the MIT Task Force on the Work of the Future, November 2020. We alter the order of presenting the six main conclusions of the MIT task force report.

[9] All of the MIT Future of Work Task Force field study reports can be found on either their Research Brief webpage https://workofthefuture.mit.edu/research-type/briefs/ or their Working Paper webpage https://workofthefuture.mit.edu/research-type/working-papers/. These two webpages also include a number of other investigative studies that were part of the overall task force effort.

[10] See the World Bank Open Data website at https://data.worldbank.org/. According to their most recent data on GDP in current US dollars, the world's 12 largest economies were US, China, Japan, Germany, India, UK, France, Italy, Brazil, Canada, Russian Republic and Korea Republic (S. Korea). Statistics on fertility rate (births per woman) and population ages 65 and above (% of total) are available through this website. The only one country of the 12 largest economies where the fertility rate was not well below replacement level was India, where it was 2.2 births per woman, and declining.

[11] Daron Acemoglu and Pascual Restrepo, "Demographics and Automation", January 2021. Forthcoming in Review of Economic Studies.

[12] Thomas H. Davenport and Steven M. Miller, "Working with Smart Machines," Asian Management Insights magazine, Vol 8 (1), May 2021, Singapore Management University. https://ink.library.smu.edu.sg/sis_research/5930/

[13] Autor, Mindell and Reynolds (2020).

[14] Erik Brynjolfsson, Seth Benzell, and Daniel Rock, "Understanding and Addressing the Modern Productivity Paradox" research brief published by the MIT Work of the Future Task Force, November 2020. A more in-depth analysis and explanation is given in Erik Brynjolfsson, Daniel Rock, and Chad Syverson, "The Productivity J-

Curve: How Intangibles Complement General Purpose Technologies," American Economic Journal: Macroeconomics, Vol 13 (1), January 2021.

[15] See Tom Davenport, "The Future of Work Now: AI-Assisted Skin Imaging," Forbes online column, November 03, 2020, https://www.forbes.com/sites/tomdavenport/2020/11/03/the-future-of-work-now-ai-assisted-skin-imaging/?sh=5606e8177e40

[16] See Tom Davenport, "The Future of Work Now: The Computer-Assisted Translator And Lilt," Forbes online column, June 29, 2020, https://www.forbes.com/sites/tomdavenport/2020/06/29/the-future-of-work-now-the-computer-assisted-translator-and-lilt/?sh=461a4e453890

[17] See David Autor, Anna Salomons, Bryan Seegmiller, "New Frontiers: The Origins and Content of New Work, 1940-2018", Working paper, MIT Economics, 26 July 2021. In their conceptual framework, augmentation is where innovations complement labor outputs, and automation is where innovations substitute for labor inputs. They examine how employment levels by occupation have changed in response to augmenting versus automating innovations in combination with market demand increases and reductions.

[18] Thomas W. Malone, Daniela Rus, Robert Laubacher, "Artificial Intelligence and the Future of Work" research brief published by the MIT Task Force on Work of the Future, December 2020.

[19] See Steven M. Miller and Thomas H. Davenport, "A Smarter Way to Manage Mass Transit in a Smart City: Rail Network Management at Singapore's Land Transport Authority," AI Singapore website, May 27, 2021, https://aisingapore.org/2021/05/a-smarter-way-to-manage-mass-transit-in-a-smart-city-rail-network-management-at-singapores-land-transport-authority/

[20] See Thomas H. Davenport and Steven M. Miller, "The Future of Work Now: The Multi-Faceted Mall Security Guard At A Multi-Faceted Jewel, " Forbes online column, September 28 2020. https://www.forbes.com/sites/tomdavenport/2020/09/28/the-future-of-work-now-the-multi-faceted-mall-security-guard-at-a-multi-faceted-jewel/?sh=2074b5ca72ff

[21] Autor, Mindell and Reynolds (2020), Section 2, Labor Markets and Growth; and Autor, Mindell and Reynolds (2019), Section 2, The Paradox of the Present, Section 3, Technology and Work: A Fraught History, and Section 4, Is This Time Different?

[22] Autor, Mindell and Reynolds (2020), Section 2.3, Employment Polarization and Diverging Job Quality.

[23] Autor, Mindell, and Reynolds (2020), Introduction.

[24] Autor, Mindell and Reynolds (2020) go on to explain in their introduction, "Four decades ago, for most U.S. workers, the trajectory of productivity growth diverged from the trajectory of wage growth. This decoupling had baleful economic and social consequences: low paid, insecure jobs held by non-college workers; low participation rates in the labor force; weak upward mobility across generations; and festering earnings and employment disparities among races that have not substantially improved in decades. While new technologies have contributed to these poor results, these outcomes were not an inevitable consequence of technological change, nor of globalization, nor of market forces. Similar pressures from digitalization and globalization affected most industrialized countries, and yet their labor markets fared better."

[25] Autor, Mindell, and Reynolds (2020), Introduction.