WHAT CAN MACHINES LEARN AND WHAT DOES IT MEAN FOR OCCUPATIONS AND THE ECONOMY?

By Erik Brynjolfsson, Tom Mitchell, and Daniel Rock

Concern about automation's impact on employment is growing as rapid advances in machine learning (ML), many based on deep neural networks, are poised to generate significant economic value and transform numerous occupations and industries. The question of whether—or to what extent--machines will replace human labor looms large today.

Our research suggests that ML technologies will indeed grow more pervasive, but within job categories, what we define as the "suitability for machine learning" (SML) of work tasks varies greatly. We further propose that our SML rubric, illustrating the variability in task-level SML, can serve as an indicator for the potential reorganization of a job or an occupation because the set of tasks that form a job can be separated and re-bundled to redefine the job. Evaluating worker activities using our rubric, in fact, has the benefit of focusing on what ML can do instead of grouping all forms of automation together.

Debates about the effects of artificial intelligence (AI) on work should shift away from the common focus on full automation of many jobs and pervasive occupational replacement, and toward the redesign of jobs and reengineering of business processes.

IN THIS BRIEF

1. Machine learning (ML) technologies will grow more pervasive.
2. Jobs are bundles of work tasks. The suitability for machine learning (SML) for work tasks varies greatly.
3. ML will rarely automate entire jobs. More often, it will lead to the reengineering of processes and the reorganization of tasks.
4. Analysis suggests that ML will affect very different parts of the workforce, including many professional jobs, compared with earlier waves of automation.
5. A shift is needed in the debate about the impact of artificial intelligence on work: Away from the focus on full automation of many jobs, and toward the redesign of jobs and processes.

MACHINE LEARNING AND THE WORKFORCE

Machine learning, a sub-field of AI, studies the question, "How can we build computer programs that automatically improve their performance at some task through experience?" Recent rapid progress in ML has made it possible for machines to match or surpass humans in certain types of tasks, especially those involving image and speech recognition, natural language processing, and predictive analytics. So far, the realized economic effects of ML are small relative to its potential. As is common, there is a time lag of years, or even decades, before technological advances generate substantial economic value: Entrepreneurs and innovators take time to adopt new technologies, co-invent complementary technologies, discover new business processes, and reconfigure existing work. This is especially true of General-Purpose Technologies (GPTs) like AI. GPTs become pervasive, improve over time, and generate complementary innovation (Bresnahan and Trajtenberg 1995).

By contrast, most of the recent progress in ML performance has been made by a specific class of algorithms called deep neural networks, or more generally, deep learning systems.

Adoption of robots, in particular, has been connected to reduced employment and wages in local labor markets (Acemoglu and Restrepo 2017). A recent study by the McKinsey Global Institute even suggested that about half of the work activities people perform could be automated with current technology (Manyika et al. 2017). While automation is already having significant effects on many parts of the workforce and advances in ML are impressive, we remain far from a world of Artificial General Intelligence (AGI) that replaces human work across entire occupations.

To delve deeper, we focused on which work tasks within occupations will be most affected by ML, and which will be relatively unaffected. When considering this question, a key insight must be maintained: An occupation can be viewed as a bundle of tasks, some of which offer better applications for technology than others (Autor, Levy, and Murnane 2003). As with other studies of task automation, the impact of ML on employment is a function of SML for specific work activities. We find that ML's potential will affect a different set of tasks than earlier technologies for task automation.

Our research examines the channels by which ML can affect the workforce. We apply Brynjolfsson and Mitchell's (2017) rubric for evaluating the potential for applying ML to 2,069 work activities, 18,156 tasks, and 964 occupations in the O*NET database. From this, we build measures of SML for labor inputs in the U.S. economy. We then discuss measures of the potential for reorganization.

WHAT CAN MACHINES LEARN AND WHAT DOES IT MEAN FOR OCCUPATIONS AND THE ECONOMY?

By Erik Brynjolfsson, Tom Mitchell, and Daniel Rock

One criterion for task SML is that the set of inputs and the corresponding set of outputs for the task can be measured sufficiently well that a machine can learn the mapping between the two sets.

In the case of ML, we find that:
1. Most occupations in most industries have at least some tasks that are SML.
2. Few occupations have all tasks that are highly SML.
3. Unleashing ML potential will require significant redesign of the task content of jobs, as SML and non-SML tasks within occupations become unbundled and re-bundled.

MACHINE LEARNING AND TASK AUTOMATION

Previous-generation automation has had a significant impact on productivity and the workforce based on explicit rules or manually written computer algorithms. However, applications of automation were limited to areas where knowledge could be codified as a computer program. Because of Polanyi's Paradox—the fact that we humans have "tacit knowledge"; we "know more than we can tell" (Polanyi 1966)— many tasks that humans know how to do, such as visually inspecting parts, had resisted automation because of our inability to codify this skill in a computer program.

By contrast, most of the recent progress in ML performance has been made by a specific class of algorithms called deep neural networks, or more generally, deep learning systems.[1] With deep learning systems, ML models circumvent Polanyi's Paradox by inferring the mapping function between inputs and outputs automatically and analyzing large amounts of sample data instead of being explicitly programmed. While not always interpretable or explainable, these ML models open a new set of possibilities for automation and complementarities to labor. Software using deep neural nets can be extended to new domains formerly closed to digitization by the high cost or impossibility of writing explicit maps of inputs to outputs and policies.

As a result, the types of tasks affected by ML tomorrow will be quite different from those affected in past waves of automation.

1 The AI Index Report at http://cdn.aiindex.org/2017-report.pdf contains a series of benchmarks.

ASSESSING SML

Successful application of ML is contingent on a variety of task characteristics and contextual factors of work activities. While we find it daunting to imagine all the ways a task could be automated—matching wits with the collective ingenuity of all the world's entrepreneurs—the scope of tasks that are SML is much more constrained and definable. Evaluating worker activities with our rubric has the benefit of focusing on what ML can do and avoiding grouping all forms of automation together.

Suboptimal bundling of SML and non-SML tasks in jobs couldalso prevent specialization and block potential productivity gains from ML. For instance, if the cost of ML capital (and SML task wage) were zero, workers would prefer to switch to tasks that ML cannot do. If firms only offer labor contracts that have a preset mixture of SML and non-SML tasks, all of the labor effort put toward SML tasks has an output opportunity cost. In other words, ML could be doing those tasks, and the firm could increase profit if it were to reorganize jobs into new bundles of tasks.

One criterion for task SML is that the set of inputs and the corresponding set of outputs for the task can be measured sufficiently well that a machine can learn the mapping between the two sets.

As noted, we used the O*NET content model for 964 occupations in the U.S. economy joined to 18,156 specific tasks at the occupation level, which are further mapped to 2,069 direct work activities (DWAs) shared across occupations. Applying a rubric of 23 questions, each DWA is scored for its SML by seven different people via the crowd-sourcing platform CrowdFlower.

The rubric is applied to each DWA to generate initial SML scores. High values of SML indicate where ML might have the greatest potential to transform a job.

	OCCUPATIONS	TASKS	DWAS
MEAN SML	3.47	3.47	3.47
STANDARD DEVIATION OF SML	0.11	0.31	0.32
MINIMUM SML	2.78	2.38	2.38
25TH PERCENTILE SML	3.40	3.25	3.25
75TH PERCENTILE SML	3.50	3.68	3.70
MAX SML	3.90	4.48	4.48
COUNT	966	19,612	2,069

Table 1-Suitability for Machine Learning: Summary Statistics

MIT INITIATIVE ON THE DIGITAL ECONOMY

WHAT CAN MACHINES LEARN AND WHAT DOES IT MEAN FOR OCCUPATIONS AND THE ECONOMY?

By Erik Brynjolfsson, Tom Mitchell, and Daniel Rock

Table 1 summarizes the SML measures for occupations, tasks, and activities from our analysis. Table 2 presents the occupations with the five highest and five lowest values for SML on a scale ranging from 1 to 5. Massage therapists seem fairly immune to machine learning technology, for instance, while concierges may be concerned. (Interestingly, the occupation "economist" scores close to average, with SML of 3.46) The variance of occupation-level SML is considerably lower than task-level SML.

TABLE 2—LOWEST AND HIGHEST 5 SML SCORE OCCUPATIONS

Low SML occupations	SML	High SML occupations	SML
Massage therapists	2.78	Concierges	3.9
Animal scientists	3.09	Mechanical drafters	3.9
Archeologists	3.11	Morticians, undertakers, and funeral directors	3.89
Public address system and other announcers	3.13	Credit authorizers	3.78
Plasterers and stucco masons	3.14	Brokerage clerks	3.78

Scale: 5 = maximum SML, 1 = minimum SML

Our research reveals a high level of variability for the potential of ML within jobs (within-occupation standard deviation of task SML scores is 0.596). Jobs with higher within-occupation standard deviation of SML have higher potential for reorganization.

There are a number of important conceptual caveats to this application of the SML rubric. First, the rubric focuses on technical feasibility. It is silent on the economic, organizational, legal, cultural, and societal factors that can have an important influence on ML adoption. Matching the evolving state-of-the-art in ML in the future will require updating the rubric accordingly.[2]

IMPLLICATIONS OF SML FOR THE WORKPLACE

It's likely that ML will not have the same effects as the last waves of automation, which led to increased inequality and wage polarization as routine cognitive tasks were automated (Autor and Dorn 2013). The correlation coefficients of SML with wage percentile and wage bill percentiles, for instance, are very low: −0.14 and 0.10, respectively. ML is a very different technology from earlier types of automation and it affects a very different set of tasks.

Furthermore, we find indicators that the next wave of automation and reengineering may affect a different part of the labor force than the last one: The correlation coefficients with wage and total wage bill percentiles and within-occupation standard deviation of SML are 0.17 and 0.002. However, it's important to note that the ex ante potential of ML may differ from its ultimate implementation, as other factors come to bear. We might see, for example, large-scale ML platform companies contracted to automate aspects of various jobs. The wage and employment effects of these contracts are ambiguous given possible channels of demand elasticity, complementary task efforts, and substitutes.

Additionally, although SML correlation with wage and total wage expenditure percentiles is low, the actual implementation of ML technologies by managers and integrators may not follow the SML rankings. If technological change is directed, the implementation of ML by managers and entrepreneurs will be focused on the high wage bill tasks with higher SML.

If ML does substitute outright for some occupational tasks, re-bundling residual tasks in new jobs may transfer risk from the firm to its workers. This will affect job design, compensation, and the organization of work. For instance, workers may need to be compensated for taking on bundles of tasks with harder-to-measure average performance when machines handle measurable tasks. Thus, over time, worker performance would become harder to evaluate since the most measurable tasks tend to be the most SML.

Improvements in technologies have historically been the key driver of increased industrial productivity. At the same time, they have also disrupted employment and the wage structure systematically. Our analysis suggests that in this era of technological progress, ML will affect very different parts of the workforce than earlier technology waves. Furthermore, tasks within jobs typically show considerable variability in SML, while few (if any) jobs can be fully automated using ML.

Machine learning technology can and will transform many jobs in the economy, but automation of entire jobs will be less significant than the reengineering of processes and the reorganization of tasks. The focus of researchers, as well as managers and entrepreneurs, should, therefore, be not just on automation, but on job redesign.

2 Rubric details are available in the Supplementary Materials to Brynjolfsson and Mitchell (2017).

WHAT CAN MACHINES LEARN AND WHAT DOES IT MEAN FOR OCCUPATIONS AND THE ECONOMY?

By Erik Brynjolfsson, Tom Mitchell, and Daniel Rock

Additional References

Acemoglu, Daron, and Pascual Restrepo. 2017. "Robots and Jobs: Evidence from US Labor Markets." Unpublished.

Autor, David H., and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review* 103(5): 1553–97.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics* 118(4): 1279–333.

Brynjolfsson, Erik, and Tom Mitchell. 2017. "What Can Machine Learning Do? Workforce Implications." *Science 358*(6370): 1530–34. Supplementary materials: http://science.sciencemag.org/content/sci/suppl/2017/12/27/358.6370.1530.DC1/aap8062-Brynjolfsson-SM.pdf

Manyika, James, Michael Chui, Mehdi Miremadi, Jacques Bughin, Katy George, Paul Willmott, and Martin Dewhurst. 2017. A Future That Works: Automation, Employment and Productivity. Chicago: McKinsey Global Institute.

Polanyi, Michael. 1966. "The Logic of Tacit Inference." *Philosophy* 41(155): 1–18

REPORT

The full paper can be found here: http://ide.mit.edu/sites/default/files/publications/pandp.20181019.pdf

ABOUT THE AUTHORS

Erik Brynjolfsson is the Director of the Initiative on the Digital Economy. He also serves as the Schussel Family Professor of Management at MIT Sloan, and Research Associate at the National Bureau of Economic Research, and is currently a Visiting Scholar at Stanford University. At MIT, he teaches courses on the Economics of Information and the Analytics Lab. He is the author of several books including, with co-author Andrew McAfee, most recently, Machine, Platform, Crowd (2017), and The Second Machine Age (2014)

Tom Mitchell is the E. Fredkin University Professor at the Machine Learning Department, School of Computer Science, at Carnegie Mellon University. His research involves the development of new machine learning algorithms, theories and applications, and studies of their societal impacts. Mitchell is a member of the U.S. National Academy of Engineering, and the American Academy of Arts and Sciences.

Daniel Rock is a PhD Candidate at MIT Sloan and a Researcher at the MIT Initiative on the Digital Economy. He studies how firms make and earn returns to investments in technology, and is particularly interested in the economics of Artificial Intelligence.

MIT INITIATIVE ON THE DIGITAL ECONOMY

The MIT IDE is solely focused on the digital economy. We conduct groundbreaking research, convene the brightest minds, promote dialogue, expand knowledge and awareness, and implement solutions that provide critical, actionable insight for people, businesses, and government. We are solving the most pressing issues of the second machine age, such as defining the future of work in this time of unprecendented disruptive digital transformation.

SUPPORT THE MIT IDE

The generous support of individuals, foundations, and corporations are critical to the success of the IDE. Their contributions fuel cutting-edge research by MIT faculty and graduate students, and enables new faculty hiring, curriculum development, events, and fellowships. Contact Christie Ko (cko@mit.edu) to learn how you or your organization can support the IDE.

TO LEARN MORE ABOUT THE IDE, INCLUDING UPCOMING EVENTS, VISIT IDE.MIT.EDU

MIT INITIATIVE ON THE DIGITAL ECONOMY