AEA Papers and Proceedings 2018, 108: 1–4
https://doi.org/10.1257/pandp.20181021

A Method to Link Advances in Artificial Intelligence to Occupational Abilities†

By Edward W. Felten, Manav Raj, and Robert Seamans*

Artificial Intelligence (AI) has great promise as an innovation that may lead to economic growth (Council of Economic Advisers 2016; Agrawal, Gans, and Goldfarb forthcoming; Brynjolfsson, Rock and Syverson forthcoming). For example, according to Graetz and Michaels (2015), robotics, an advanced technology with similarities to AI, added an estimated 0.37 percentage points of annual GDP growth between 1993 and 2007 on average for the 17 countries in their sample (accounting for about one-tenth of GDP growth during this time period). The authors note that these effects are of similar magnitude to the impact of steam engines on growth in the United Kingdom.

However, while AI may boost growth, the effect on labor is less clear. Historically there is empirical evidence that automation can both complement and substitute for labor (Autor and Salomons 2017; Bessen 2017). In the specific case of robots, research provides mixed findings, with some researchers finding no effect of robots on labor (Graetz and Michaels 2015), and others finding evidence that robot adoption leads to job losses (Acemoglu and Restrepo 2017). To date, however, there has been little systematic empirical research on the link between AI and labor, and the handful of existing studies arrive at different findings

In a widely cited paper, Frey and Osborne (2017) categorize tasks by their susceptibility to automation, link these tasks to occupation, employment, and wage data, and find that 47 percent of US employment is at high risk of automation. One assumption embedded in the Frey and Osborne model is that all workers in the same occupational category face the same threat of automation. An OECD Report (Arntz, Gregory, and Zierahn 2016) instead argues that there may be task variation between individuals within the same occupation. For example, managers of different firms may treat shop-floor labor differently, depending on whether they view workers as partners in the production process or as inputs into a production function (Helper, Martins, and Seamans 2018). The OECD Report instead uses individual level data to predict how susceptible occupations may be to automation, and finds that only 9 percent of jobs in the United States and across OECD countries will be highly susceptible to automation.

Our paper provides a new method that we believe can help researchers and policymakers to better understand the link between AI and labor. We follow recent work in economics that describes "labor" via the bundle of skills or abilities that are used for any specific occupation (e.g., Autor and Handel 2013; Brynjolfsson, Mitchell, and Rock 2018). Our method—which is described in detail below—links advancement in different categories of AI to different types of abilities. The effect of advancement in AI on abilities can then be aggregated to occupations and industries. Our approach complements that of Frey and Osborne (2017) by relying on third party measures of past advances in AI rather than on experts' predictions of the future, and complements that of Brynjolfsson, Mitchell, and Rock (2018) by estimating how AI has advanced over time. In principle, our approach allows other researchers, practitioners, and policymakers to model how advances in AI affect different abilities, occupations, and industries. We also provide a test of our method that links advancement in AI

* Felten: Princeton University, Sherrerd Hall, Third Floor, Princeton, NJ 08544 (email: ed@felten.com); Raj: NYU Stern School of Business, 44 West 4th Street, New York, NY 10012 (email: mraj@stern.nyu.edu); Seamans: NYU Stern School of Business, 44 West 4th Street, New York, NY 10012 (email: rseamans@stern.nyu.edu). We thank Susan Helper and AEA participants for helpful comments, discussions, and feedback.

† Go to https://doi.org/10.1257/pandp.20181021 to visit the article page for additional materials and author disclosure statement(s).

categories between 2010 and 2015 to subsequent updates in occupational descriptions.

I. Linking Advances in AI to Abilities

Our method relies upon two independent databases—the Electronic Frontier Foundation (EFF) AI Progress Measurement dataset and the Occupational Information Network (O*NET) database developed by the US Department of Labor.[1]

The EFF AI Progress Measurement experiment is a pilot project that aims to track progress on task-specific AI performance metrics across a variety of separate artificial intelligence categories, such as abstract strategy games and image recognition, for example. For each of the categories, the EFF monitors progress in the field drawing on data from a variety of sources, including blog posts and websites focused on subfields of machine learning, academic literature, and review articles. The EFF aims to create the first integrated database that aggregates performance metrics of state of the art systems across a variety of artificial intelligence categories in one single place, and therefore to provide researchers, policymakers, and technology users with insight into the state and the rate of development of the field.

The O*NET database is a comprehensive database that provides occupational definitions for professions in the modern day American workplace. Since the 1990s, the US Department of Labor has developed and maintained the database to provide up-to-date information as the nature of the occupations listed changes. For each of the almost 1,000 occupations listed, O*NET provides information regarding personal requirements, personal characteristics, experience requirements, job requirements, and the state of the labor market. For the purposes of our study, we focus on job requirements. O*NET maintains a list of 52 distinct abilities, and in each occupation's job requirements, it notes how important and prevalent each ability is in the relevant occupation.

We use the EFF AI Progress Measurement dataset to track the rate of change across the 16 separate categories of metrics the EFF tracks. For each of the categories, we first integrate all the different metrics tracked to get a comprehensive understanding of the pace of progress in the AI subfield corresponding to the category of metrics. This can be an intricate process, as measures within a category can utilize different scales and present distinct results. To provide an illustrative example, Figure 1 shows the data for the various metrics of image recognition tracked by the EFF.

For the image recognition category, the EFF provides seven separate metrics. To calculate the slope measuring the progress in image recognition as a whole, each metric must first be scaled appropriately. For example, if the metric is error rate on some task, we scale by taking the negative logarithm of the error rate, yielding a scaled metric that will grow linearly if the error rate is decreasing exponentially. Next, we fit a model which assumes a single linear rate of increase in the scaled metrics, plus a per-metric offset. The rate of increase found by this method serves as our estimate of the progress rate for image recognition. For some AI categories, at the time of publication, the EFF either provided very little or no information regarding past progress. For those categories, the slope measuring progress was set equal to zero.

Next, we map the EFF AI categories to the list of 52 abilities that the O*NET database uses to describe job requirements. To do so, we construct a matrix that connects the two. The matrix was constructed using inputs from multiple computer science PhD students. With the matrix, we are able to connect the EFF categories to the O*NET abilities, and can then measure the relative effect of advances in AI technology on the different abilities listed by O*NET. We can then use the O*NET occupational definitions to evaluate the impact of AI technology advances on each occupation by weighting the effect of AI technology on each ability by the ability's prevalence and importance for each job. We aggregate the impact across all abilities at the occupation-level to create an effect score for each occupation. While the value of the score itself is arbitrary, it allows us to compare the relative impact of AI technology *across* a variety of occupations.

[1] AI Progress Measurement from Electronic Frontier Foundation is available at https://www.eff.org/ai/metrics. A description of the Department of Labor's O*NET abilities is available at https://www.onetonline.org/find/descriptor/browse/Abilities/.

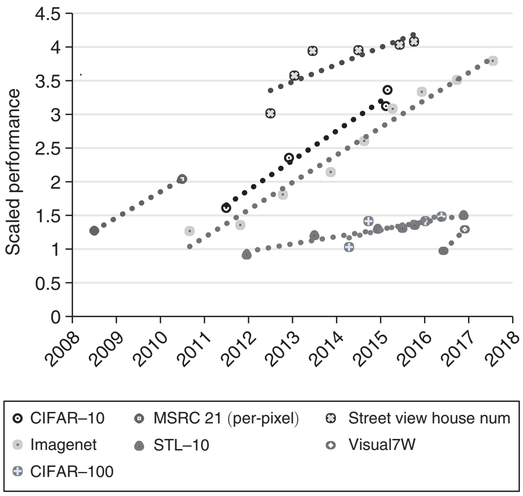

Figure 1. Progress across EFF Image Recognition Measures

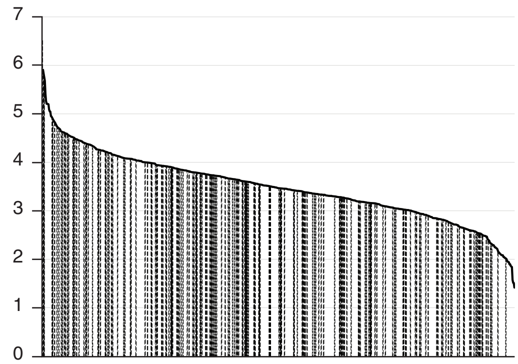

Figure 2. Distribution of Occupation Employment Effect Scores Based on Historical Progress of EFF Metrics

II. Historical Progress in AI

Using the progress slopes as calculated above, we were able to identify a list of occupations that were the most and least impacted by AI technology over the last few years. To check the validity of our methodology, we examined the correlation between the occupation-level impact score and whether the BLS was planning on changing the official occupational definition for each job in 2018. The last updates to BLS occupation definitions were in 2010, so presumably, the occupations most impacted by AI between 2010 and 2015, when decisions were made regarding which occupations to update, would be more likely to have changed in nature and require an update of their BLS definition.

Figure 2 graphically charts the distribution of the occupational impact scores, where the approximately 1,000 occupations are ordered from most to least affected by advancements in AI between 2010 and 2015. The solid black line shows the overall distribution of the occupational impact scores. The dashed columns represent occupations that will be receiving updated definitions in 2018. One-hundred-five of the updated occupations are above the median occupational impact score, and 83 are below the median.

We conducted an analysis to identify whether there was any statistically significant correlation between an occupational impact score and whether an occupation was scheduled to receive a definition change. We found a statistically significant correlation coefficient of 0.074 $(p = 0.041)$ between the impact score and a scheduled definition change. Because the impact score is arbitrary, it is difficult to interpret the magnitude of this coefficient, however, it confirms a positive and significant relationship between the impact scores and definition changes. Of course, many other factors, including new product and process innovations and international trade, likely also affect whether and how an occupation changes over time.

Finally, note that our methodology does not speak to whether AI is serving as a substitute or complement to the occupations it effects—rather, it only suggests which occupations require abilities that may be affected by advances in AI technology, and we believe these effects can be either substitutes or complements. For example, researchers and policymakers could use our method to identify which occupations will be most (or least) affected by a simulated 10 percent advancement in the application of AI to image recognition, speech recognition, or other AI progress category. Policymakers and researchers could then use the occupation lists generated from such simulations to focus on a narrower set of occupations for further study and research.

III. Implications for Future Work

The big question that has grabbed policymakers and pundits is will artificial intelligence take all the jobs? In order to understand the effects of AI on labor, however, more work needs to be done linking advances in AI to occupations and skills.

In this paper, we develop such a methodology, and use it to correlate advances in AI to actual changes to occupational descriptions. Our methodology should be useful to other researchers, practitioners, and policymakers studying the effect of advances in AI on skills, occupations, and industry. For example, future studies could make use of our methodology to study how a rapid increase in certain types of AI may have distributional effects that vary by occupations, industry, or geography. Our methodology would benefit from more research to create a more systematic link between AI categories and abilities.

REFERENCES

Acemoglu, Daron, and Pascual Restrepo. 2017. "Robots and Jobs: Evidence from US Labor Markets." National Bureau of Economic Research Working Paper 23285.

Agrawal, Ajay, Joshua S. Gans, and Avi Goldfarb. Forthcoming. "Prediction, Judgment, and Complexity." In *Economics of Artificial Intelligence***,** edited by Ajay K. Agrawal, Joshua S. Gans, and Avi Goldfarb. Chicago: Univeristy of Chicago Press.

Arntz, Melanie, Terry Gregory, and Ulrich Zierahn. 2016. "The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis." Organization for Economic Cooperation and Development Working Paper 189.

Autor, David H., and Michael J. Handel. 2013. "Putting Tasks to the Test: Human Capital, Job Tasks, and Wages." *Journal of Labor Economics* 31 (2): S59–96.

Autor, David H., and Anna Salomons. 2017. "Robocalypse Now—Does Productivity Growth Threaten Employment?" Unpublished.

Bessen, James E. 2017. "Automation and Jobs: When Technology Boosts Employment." Boston University School of Law Research Paper 17–09.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. 2018. "What Can Machines Learn, and What Does it Mean for the Occupations and Industries?" Unpublished.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. Forthcoming. "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics." In *Economics of Artificial Intelligence*, edited by Ajay Agrawal, Joshua S. Gans, and Avi Goldfarb. Chicago: University of Chicago Press.

Council of Economic Advisers. 2016. *2016 Economic Report of the President*. Washington, DC: Council of Economic Advisers.

Frey, Carl Benedikt, and Michael A. Osborne. 2017. "The Future of Employment: How Susceptible are Jobs to Computerisation?" *Technological Forecasting and Social Change* 114: 254–80.

Graetz, Georg, and Guy Michaels. 2015. "Robots at Work." Centre for Economic Performance Discussion Paper 10477.

Helper, Susan, Raphael Martins, Robert Seamans. 2018. "Value Migration and Industry 4.0: Theory, Field Evidence, and Propositions." Unpublished.