JOB AUTOMATION RISK AND THE FUTURE OF SKILLS: SKILLS AND COMPETENCY CHANGE IN THE U.S. WORKFORCE

Chief Evaluation Office Evaluation Technical Support Project
Contract # 1605DC-18-A-0019/1605DC-F-00431

May 15, 2023

Submitted to:
Monica Mean
Sr. Evaluation Specialist
Chief Evaluation Office
U.S. Department of Labor
200 Constitution Ave NW, Room
Washington, DC 20210

Submitted by:
Marilia G. Mochel
Project Director
Manhattan Strategy Group
4340 East-West Hwy Suite 1100
Bethesda, MD 20817

Manhattan Strategy Group
4340 East-West Highway, Suite 1100
Bethesda, MD 20814
P: 301.828.1515
E: info@manhattanstrategy.com
W: manhattanstrategy.com

This report was prepared for the U.S. Department of Labor (DOL), Chief Evaluation Office (CEO), by Manhattan Strategy Group under contract number 1605DC-18-A-0019. The views expressed are those of the authors and should not be attributed to DOL, nor does mention of trade names, commercial products, or organizations imply endorsement of same by the U.S. Government.

TABLE OF CONTENTS

INTRODUCTION

In March 2019, the U.S. Government Accountability Office (GAO) recommended that "the Secretary of Labor should direct the Bureau of Labor Statistics (BLS) and the Employment and Training Administration (ETA) to develop ways to use new and existing data collection efforts to identify and systematically track the workforce effects of advanced technologies" (Government Accountability Office, 2019, p. 2).

GAO's report came as advances in these technologies are changing the workplace, and the potential impacts of these changes could be significant to American workers. The report noted that existing workforce data that has been regularly collected by the U.S. Department of Labor (DOL) may not be able to identify the causes of employment shifts, including whether shifts are due to the dissemination of these technologies in the workplace. The GAO report underscored that knowledge of how new technologies affect employment shifts is critical for DOL's mission to fund programs that support workers. Specifically, GAO noted that the Occupational Information Network (O*NET) collects valuable data on tasks and technologies across occupations, but the system was not designed to track changes over time. GAO recommendations noted possibilities to address these gaps, including action by BLS and the O*NET program within ETA (Government Accountability Office, 2019).

At the same time, there is growing recognition that it is no longer possible to capture skill and competency levels solely through surveys and focus groups; real-time data from online and administrative databases also needs to be leveraged. In the past decade, advances in data analytics may allow for new capabilities in understanding the potential effects of the adoption of automation and artificial intelligence[1] (AI) skill and competency requirements in the workplace. DOL would benefit from learning about feasible and optimal ways to assess skill and competency change over time, as the Department is tasked to foster, promote, and develop the welfare of wage earners and jobseekers in the United States (U.S. Department of Labor, n.d.).

ETA and the Chief Evaluation Office (CEO) are planning to convene a roundtable of experts to discuss these data needs and approaches in the fall of 2022. The planned research roundtable will gather a small number of experts exploring competencies and skills change over time in their research to share their perspectives on ways to improve data available to researchers in this field. The four overarching goals for this research roundtable pertaining to the future of work are to:

1. Learn from researchers in the field various perspectives about the feasibility of analyzing skill and competency changes over time.
2. Explore methods or approaches, such as AI and machine learning (ML), to assess skills and competency demands and trends over time, as certain occupations increase or decrease in size or as task composition of occupations change.
3. Explore needed data sources (e.g., job postings or worker/employer surveys) to assess skill and competency changes for occupations.

[1] Muro et al. (2019) define AI as consisting of "a diverse set of technologies that serve a variety of purposes. […] Broadly speaking, AI involves programming computers to do things which—if done by humans—would be said to require 'intelligence,' whether it be planning, learning, reasoning, problem-solving, perception, or prediction."

4. Identify options to support such analytical capabilities either through government or public-private partnerships to build capacity for capturing and assessing such information (e.g., federally funded research and development centers, university partnerships).

Manhattan Strategy Group (MSG) is supporting ETA and CEO in this endeavor. As a first step, we present this scan of the literature on technological advances and U.S. jobs. This literature scan covers select research that examines how a new wave of technological innovations in computers, robotics, AI, and ML may change the types of jobs Americans get, where they get them, and how well they are paid. This review started with 12 papers provided by ETA and expanded as the bibliography list grew. We did not conduct a search for terms on academic databases but rather followed up in snowball fashion after select literature cited in the initial list of 12 papers. As a result, this literature review is by no means comprehensive, but it provides an overview of the research focused on technological changes and the U.S. labor market as well as the research approaches that generated these findings.

This report starts with a discussion of the literature on how researchers assess automation risk and (potential) impacts on the labor market. It describes methodologies used, emphasizing how these methodologies make important distinctions in the types of technological advances. The literature scan also describes the use of data sources in this field of research along with approaches to explore skills and skills change. We additionally describe how researchers recognize ambiguity in predicting the impact of automation. We close by reviewing recommendations related to data sources and methods, including O*NET, to improve our understanding related to skill changes in the labor market.

UNDERSTANDING JOB AUTOMATION RISK IN THE LABOR MARKET

A growing body of research has been published that builds on the insight that technology's influence on the labor market depends on the capabilities new technologies bring to the workplace and on the composition of the tasks workers perform across different jobs. As articulated early in Autor et al. (2003), writing specifically on the adoption of computers in the workplace, the fit between activities or tasks new technologies can perform and what jobs consist of has implications on whether a specific job is likely to be replaced or complemented by the adoption of the technology. These two elements—the view of jobs as a combination of tasks and the different suitability of different tasks to computerization or other technologies—are at the core of many studies in this field, with variations depending on the technology being studied.

This literature scan therefore examines research that uses a variety of data sources to explore (1) the possible uses for new technology in the workplace, (2) the composition of tasks workers perform across occupations and industries, (3) trends in occupations in the economy and across industries and geographical areas, and (4) the dissemination of technology in the workplace.

In the case of possible uses of new technologies in the workplace, studies use data sources such as expert assessments of technology applications (Frey & Osborne, 2017) or descriptions of applications as they appear in new technology patent applications (Webb, 2019) to assess the applicability of technologies to job tasks. To identify what types of activities individuals perform at work, studies use sources such as government databases on occupations (such as O*NET in the United States), surveys (such as Job Openings and Labor Turnover Survey, or JOLTS), and job descriptions gathered in online vacancy listings (Acemoglu et al., 2022). Surveys are also typically used to assess trends in occupations, including showing breakdowns across industries and areas. Finally, researchers have explored the dissemination of computers and robots using industry sales figures (in the case of robots) and whether the use of computers is listed in job descriptions.

Our scan of the literature describes research in this field, with attention to new measures developed for assessing automation risks and data sources used. We start by describing the foundational insight of Autor et al. (2003) regarding computerization in the U.S. workplace.

Autor et al. (2003) theorize about how the use of computers ("computerization") changes the demand for workers' skills. The authors sought to explain the roots of the well-established association between adoption of computers and the rise of college-educated workers across industries, firms, and plants. To do so, they organize job tasks across occupations into four categories: routine, nonroutine, analytic and interactive, and manual tasks. They argue that, depending on their type, tasks are more or less susceptible to replacement by computers. For tasks that can be accomplished by following explicit rules, computers can substitute for workers. Computers will instead complement worker activities when workers perform nonroutine problem-solving and complex communications tasks. The predictions of their model along with examples of tasks in each category appear in Exhibit 1, reproduced from their study.

Exhibit 1: Task Types and the Predicted Impact of Computerization.

	Routine Tasks	Nonroutine Tasks
Analytic and Interactive Tasks	Substantial substitution (Example: record-keeping, repetitive customer service)	Strong complementarity (Example: forming and testing hypotheses, legal writing, persuading/selling, and managing others)
Manual Tasks	Substantial substitution (Example: picking or sorting, repetitive assembly)	Limited opportunities for substitution or complementarity (Example: janitorial services, truck driving)

Note. From "The Skill Content of Recent Technological Change: An Empirical Exploration," by D. H. Autor, F. Levy, and R. J. Murnane, 2003, *The Quarterly Journal of Economics, 118*(4), p. 1286 (https://doi.org/10.1162/003355303322552801).

Autor et al. (2003) used the Dictionary of Occupational Titles, which preceded O*NET, combined with worker survey data from the Census and the BLS Current Population Survey, to examine changes in tasks within industries, education groups, and occupations. Their findings show that starting in the 1970s, routine cognitive and manual labor tasks declined while nonroutine analytic and interactive tasks rose. These shifts are concentrated in industries experiencing rapid computerization and affect all educational groups. Occupations experiencing rapid rise of computers also see this pattern, with reduced input of routine cognitive tasks and increased input of nonroutine cognitive tasks.

Exhibit 2, reproduced from the paper, shows these trends. It shows that, over the period of 1960 to 1998, changes in occupational distribution resulted in changes in the tasks performed by the U.S. labor force.

Exhibit 2: Trends in Routine and Nonroutine Task Input, 1960 to 1998

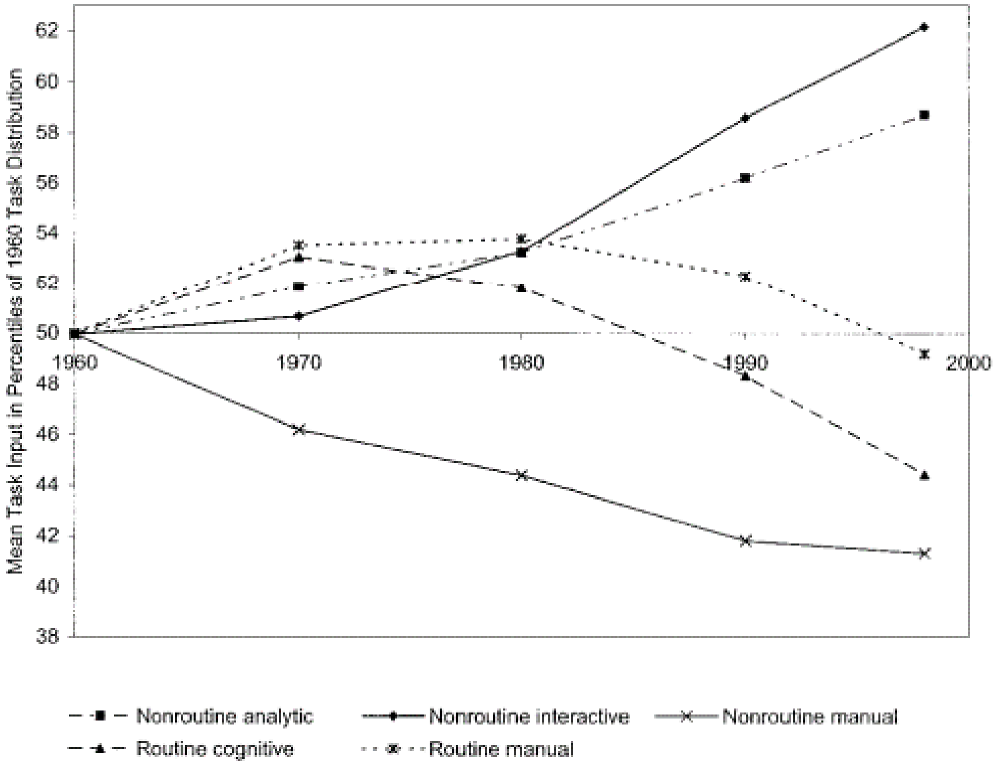

Note from the authors: Figure I is constructed using Dictionary of Occupational Titles [1977] task measures by gender and occupation paired to employment data for 1960 and 1970 Census and 1980, 1990, and 1998 Current Population Survey (CPS) samples. Data are aggregated to 1120 industry-gender-education cells by year, and each cell is assigned a value corresponding to its rank in the 1960 distribution of task input (calculated across the 1120, 1960 task cells). Plotted values depict the employment-weighted mean of each assigned percentile in the indicated year.

Note. From "The Skill Content of Recent Technological Change: An Empirical Exploration," by D. H. Autor, F. Levy, and R. J. Murnane, 2003, *The Quarterly Journal of Economics, 118*(4), p. 1296 (https://doi.org/10.1162/003355303322552801).

AUTOMATION RISK MEASURES

The past decade has seen a crop of new indicators seeking to measure risk of the dissemination of new technologies across occupations. The measures are designed to focus on specific technologies. They use various data sources to estimate the replaceability of elements of an occupation (or job) by machines or technologies.

Autor and Dorn (2009) assess how shifts in occupational structure have affected the job composition of young and older workers at different education levels between 1980 and 2005. They argue that occupational change over time sheds light on opportunities faced by workers at different ages and education levels. They explain the dynamics as follows:

> "When an occupation declines, […] older workers will face an incentive not to exit the occupation while younger workers will face an incentive not to enter. Moreover, firms may react to changing demands for occupations by hiring young workers into growing occupations and curtailing such hiring into contracting jobs. These suppositions imply that occupations will 'get old' as their employment declines–that is, the mean age of an occupation's workforce will rise." (Autor & Dorn, 2009, p. 2)

To assess their hypotheses, they use occupation level and job data on task requirements associated with the occupations (manual, routine, and abstract) from the fourth edition of DOL's Dictionary of Occupational Titles and their corresponding Census occupational classifications. They use these sources to compute a routine task intensity index (RTI), which by design rises or falls depending on the relative importance of manual tasks for the occupation. In their analysis, the authors then overlay RTI to geographical regions (commuting zones, or CZs). CZs were created with confidential commuting data from the 1990 Census to identify commuting ties within and across clusters of counties. The authors explore how shifts in the RTI in these areas are associated with employment and wages using Census employment data. They describe their findings as follows:

> "CZs that were initially specialized in routine task-intensive occupations saw substantial declines in the share of workers employed in these occupations between 1980 and 2005. CZs that were initially specialized in routine task-intensive occupations saw substantial declines in the share of workers employed in these occupations between 1980 and 2005. Relative declines in routine occupation employment within CZs are primarily offset by relative employment gains in low-skill nonroutine occupations—jobs that are significantly less skill-intensive and lower-paying than the routine occupations that are displaced." (Autor & Dorn, 2009, p. 49)

As a result, Autor and Dorn (2009, 2013) argue that workplace computerization is a key driver of rising polarization in employment and wages in the United States. They find that routine-intensive occupations saw declines in employment shares between 1980 and 2005. They also see rising wages in the period. They suggest that this pattern, where traditional clerical tasks are replaced by automation, leaves the remaining work content to be concentrated in more skill-intensive tasks. As a result, CZs with a high prevalence of routine-intensive occupations develop high job-market polarization. Autor and Dorn (2013) further show that this is due to a rapid rise of employment in service occupations, which is more pronounced at first in CZs with highly routine task-intensive labor markets. This also appears as a pattern, confirmed in their analysis, of rising employment and earnings of non-college workers in nonroutine-intensive occupations and declining employment and wages in routine-intensive occupations.

In a highly influential paper, Frey and Osborne (2013) propose a methodology to estimate the share of jobs that are susceptible to computerization. To do so, they follow Autor et al.'s (2003) two-by-two matrix, showing routine/nonroutine tasks in one axis and cognitive/manual tasks on the other. With knowledge of recent advances in technology, the authors describe engineering bottlenecks that prevent task automation at our current technological stage. These are areas in technological development that involve perception and manipulation, creative intelligence, and social intelligence.

The authors worked with a group of ML researchers at the Oxford University Engineering Sciences Department to subjectively label 70 occupations from O*NET as either automatable or not. The answers by these experts were used to model and score the remainder of the occupations on risk of automation using ML techniques. They used O*NET information to rank occupations according to the mix of knowledge, skills, and abilities they require and to categorize them based on the variety of tasks they involve. They linked O*NET occupations thus characterized to the 2010 Standard Occupational Classification system and BLS employment and wage data. This yielded a dataset comprising 702 occupations with O*NET's standardized key features and open-ended descriptions of tasks for each occupation. Occupations with a probability of automation of

70% or higher are considered high-risk, and the authors estimate that 47% of U.S. employment is at high risk of automation in the next two decades.

Unsurprisingly, Frey and Osborne's (2013) estimates made headlines and raised alarm.[2] The authors were nonetheless careful to note that these estimates do not mean to predict where automation will happen as many factors influence the decision to automate.

Arntz et al. (2016) argue that Frey and Osborne's (2013) model may overestimate the risk of automation for occupations. They maintain that its assessment of whole occupations fails to examine job tasks at a more granular level. When they do so, the authors find that occupations labeled as high risk in the Frey and Osborne study include tasks that are hard to automate. This mix of automatable and non-automatable tasks suggests that certain jobs are less likely to be automated. For example, Frey and Osborne estimate that "Bookkeeping, Accounting, and Auditing Clerks" to have a 98% automation risk, whereas Arntz et al. find that only 24% of those employed in these occupations report they can perform their job without group work or face-to-face interactions, the types of tasks that are less susceptible to automation. Accordingly, the prevalence of processes and activities would mean these jobs are not high risk.

To address this limitation, Arntz et al. (2016) used individual-level survey data that identified job tasks. Data from the Programme for the International Assessment of Adult Competencies[3] (PIAAC) provided them with details on skills, jobs, tasks, and competencies. The authors categorize job activities as more or less susceptible to automation. They find that jobs that "require cooperation with other employees or where people spend more time on influencing others" have a lower risk of automation, whereas automatability is higher in jobs "with a high share of tasks that are related to exchanging information, selling, or using fingers and hands" (p. 14).

Arntz et al. (2016) use this measure to examine the susceptibility to computerization of occupations across 21 Organization for Economic Cooperation and Development (OECD) countries. Using the individual-level assessment of job tasks, they find that, on average, 9% of jobs are highly automatable (defined as an automatability of at least 70%, as in Frey and Osborne, 2017), with large variations across the OECD countries. The estimates for the U.S. using Arntz et al.'s task-based measure is 9%. This estimation is much lower than Frey and Osborne's calculation of 47% of U.S. jobs being at risk for automation. The difference is that many occupations expected to be at high risk, when more closely examined at the job-task level, reveal that workers are actually performing activities that are difficult to automate.

Brynjolfsson and Mitchell (2017) focus more narrowly on the potential impact of advances in ML, a subset of AI applications. For a task to be performed by ML, Brynjolfsson and researchers in another study argue that:

---

[2] Examples of articles on Frey and Osborne's 2013 findings: https://www.economist.com/graphic-detail/2018/04/24/a-study-finds-nearly-half-of-jobs-are-vulnerable-to-automation

https://slate.com/technology/2013/09/researchers-claim-many-jobs-at-risk-for-automation-here-s-what-they-missed.html

https://www.nytimes.com/2016/02/28/magazine/the-robots-are-coming-for-wall-street.html

[3] PIAAC is a program of assessment and analysis of adult skills. The major survey conducted as part of PIAAC is the Survey of Adult Skills to measure adults' proficiency in key information-processing skills—literacy, numeracy, and problem-solving. The survey collects information and data on how adults use their skills at home, at work, and in the wider community. This international survey is conducted in over 40 countries/economies.

> "the set of actions 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." (Brynjolfsson et al., 2018a, pp. 44–45)

Brynjolfsson and Mitchell (2017) maintain that ML deep-learning approaches have been able to "match or surpass humans in certain types of tasks, especially those involving image and speech recognition, natural language processing, and predictive analytics" (p. 43). This has applications to many jobs, but based on the nature of ML, the authors expect the occupations affected by ML advances will be different from those occupations affected by previous waves of automation (e.g., robotics, computers) as they will impact professional jobs (e.g., credit authorizers).

Brynjolfsson and Mitchell (2017) developed a rubric to measure the exposure of occupations to these new capabilities called Suitability for Machine Learning (SML). They extracted the O*NET content for 964 occupations and 18,156 tasks at the occupation level. The tasks are mapped to 2,069 O*NET detailed work activities shared across the occupations—they used coders hired via a crowdsourcing platform to apply their 23-question rubric to provide an SML score for each. By computing the SML scores for occupations, tasks, and detailed work activities, they find variation in potential ML effects across occupations. The authors predict that "ML will affect very different parts of the workforce than earlier waves of automation" (p. 44). Job functions that involve classification (e.g., labeling of medical records) and predication (e.g., analyzing loan applications) are some of the applications the authors find most suitable to ML applications. They also note that most occupations include some SML tasks, but few are fully automatable. As a result, they expect that many occupations are likely to be reorganized instead of replaced.

Focusing on advances in AI more broadly, Felten et al. (2019) developed the Artificial Intelligence Occupational Impact (AIOI) measure. The AIOI is built using the Electronic Frontier Foundation (EFF) AI Progress Measurement dataset and O*NET data. The EFF AI Progress Measurement project[4] "tracks reported progress on metrics of AI performance across separate artificial intelligence applications, such as image recognition, speech recognition, translation, or abstract strategy games, drawing on data from multiple sources, including academic literature, review articles, blog posts, and websites focused on artificial intelligence" (Felten et al., 2019, p. 2). Felten et al. create AIOI that measures how advances in AI from 2010 to 2015 (period covered by EFF data) used 52 abilities descriptors rated in O*NET. These abilities, unlike activities, are "designed to capture something more fundamental about what an individual brings to a given occupation" (p. 8). The authors use a crowdsourcing platform to connect the abilities and AI progress data using results of a survey of 1,800 participants. The impact of AI across abilities is then weighed by the sum of the prevalence and importance of all abilities in the occupation.

Felten et al. (2019) use this measure in combination with BLS data for each occupation at the state level from 2010 to 2016 to assess employment and wages at the state-occupation level and to explore whether AI replaces or complements work in these occupations. Their analysis, which incorporates data on the prevalence of software skills within an occupation in 2010 from Burning

[4] For more information on the EFF AI Progress Measurement project, see https://www.eff.org/ai/metrics.

Glass,[5] identifies that the positive relationship between AI and wage growth is driven by occupations involving a high level of software skills. The analysis also shows that AI may have positive impacts on higher-income occupations, with no significant impact on low- and middle-income occupations.[6]

Felten et al. (2021) expand and refine the measure, now called the AI Occupational Exposure (AIOE), by scaling the aggregate exposure to AI at the occupation level by the abilities used in each occupation to better measure exposure at the occupation level. Their measure focuses on 10 of the fastest-growing AI applications. They also expand the applications of the AIOE to the industry level by taking a weighted average of the AIOE using industry employment based on the four-digit North American Industry Classification System classification in 2019. Exhibit 3 combines two tables from Felten et al. (2021). They estimate the highest and lowest AIOE scoring occupations and the industries with highest and lowest scores in AI Industry Exposure (AIIE), a measure of AI exposure at the industry level based on AIOE. The AIOE measure indicates that the AI dissemination will affect occupations with tasks that require problem-solving, logical reasoning, and perception as opposed to physical capabilities. Exhibit 3 combines in one table Felten et al.'s 10 highest and lowest AIOE scoring occupations and the industries with highest and lowest scores in AIIE.

Exhibit 3: Occupations and Industries with the Highest and Lowest AIOE/AIIE Measures

Rank	Hghest-scoring occupation	Lowest-scoring occupation	Highest-scoring industry	Lowest-scoring industry
1	Genetic counselors	Dancers	Securities, commodity contracts, and other financial investments and related activities	Support activities for crop production
2	Financial examiners	Fitness trainers and aerobics instructors	Accounting, tax preparation, bookkeeping, and payroll services	Services to buildings and dwellings
3	Actuaries	Helpers—painters, paperhangers, plasterers, and stucco masons	Insurance and employee benefit funds	Foundation, structure, and building exterior contractors
4	Purchasing agents, except wholesale, retail, and farm products	Reinforcing iron and rebar workers	Legal services	Animal slaughtering and processing
5	Budget analysts	Pressers, textile, garment, and related materials	Agencies, brokerages, and other insurance related activities	Building finishing contractors

[5] Burning Glass data gathers and integrates economic, labor market, demographic, education, profile, and job posting data from dozens of government and private-sector sources, creating a comprehensive and current dataset that includes both published data and detailed estimates with full United States coverage.

[6] The authors use median annual income for an occupation in 2010, splitting the sample in terciles to identify the three levels of income.

Rank	Hghest-scoring occupation	Lowest-scoring occupation	Highest-scoring industry	Lowest-scoring industry
6	Judges, magistrate judges, and magistrates	Helpers—Brickmasons, Blockmasons, stonemasons, and tile and marble setters	Nondepository credit intermediation	Warehousing and storage
7	Procurement clerks	Dining room and cafeteria attendants and bartender helpers	Other investment pools and funds	Fiber, yarn, and thread Mills
8	Accountants and auditors	Fence erectors	Insurance carriers	Support activities for rail transportation
9	Mathematicians	Helpers—roofers	Software publishers	Sawmills and wood preservation
10	Judicial law clerks	Slaughterers and meat packers	Lessors of nonfinancial intangible assets (except copyrighted works)	Support activities for water transportation

Note from the authors: Occupations are ranked by their constructed AIOE measure at the six-digit Standard Occupational Classification level. Occupation titles are taken from the O*NET database. Highest-scoring occupations are ranked in descending order based on the AIOE measure. Lowest-scoring occupations are ranked in ascending order based on the AIOE measure. Industries are ranked by their constructed AIIE measure at the four-digit North American Industry Classification System level. Industry titles are taken from the BLS. Highest-scoring occupations are ranked in descending order based on the AIIE measure. Lowest-scoring occupations are ranked in ascending order based on the AIIE measure (Felten et al., 2021).

Webb (2019) also focuses on how developments in AI technology may affect occupations in the United States. He proposes a new methodology to measure exposure of occupations to AI that uses patent data. His measure captures the overlap between the text of patents and job applications and occupation descriptions. He uses Google Patents Public Data, including title, abstract, and Cooperative Patent Classification codes, and the text in O*NET, including the importance and frequency of each task within each occupation. These two sources can be combined to quantify how much patenting in a particular technology has been directed at the tasks of any occupation. This is a measure of the tasks from which labor may be displaced.

To test this new measure, Webb (2019) provides historical empirical analysis. Using Census employment data from 1960–2000 and American Community Survey data from 2000–2018, he finds that occupations with high exposure to previous automation technologies (industrial robots and software development) see declines in employment and wages. His analysis also shows differences in the impact of these two technologies. On the one hand, individuals with less than a high school education in low-wage occupations are more exposed to robots. On the other, those in middle-wage occupations are the most exposed to software technologies. In his analysis, Webb then looks prospectively to the effects of AI exposure. He finds a different worker profile exposure when it comes to AI applications in the workplace: the highest exposure group in his measure appears to be those in highly skilled occupations. Webb finds that AI is more suitable to "tasks that involve detecting patterns, making judgments, and optimization. Most-exposed occupations include clinical laboratory technicians, chemical engineers, optometrists, and power plant operators" (Webb, 2020, p.3)

Mann and Püttmann (2021) also use data on all U.S. patents granted between 1976 and 2014 to examine the effects of automation on U.S. labor markets at the CZ level. They use ML to distinguish between automation and non-automation innovations based on the text of patents granted between 1976 and 2014. They then probabilistically assign automation patents to specific industries and create an analytic dataset of U.S. CZs with this industry-level data, covering 722 CZs over 39 years. The authors find that automation is beneficial to local employment, which suggests that the worker substitution effect of automation, whereby technology would lead to layoffs in certain occupations, is more than compensated by growth in product demand corresponding to growth in other occupations.

The literature on automation's impact on workers has also examined robotics and how the use of robots in industry affects workers. Graetz and Michaels (2018) find that robot adoption predicts wage growth and lower consumer prices, but employment then shifts from low-skilled workers to middle- and higher-skilled workers. They used data from the International Federation of Robotics (IFR) on deliveries of industrial robots to estimate "robot densification" (the stock of robots per million hours worked) across 14 industries in 17 countries from 1993 to 2007. They also used the IFR data to classify the tasks performed by robots (robot applications). Applications were matched to the 1980 occupational types to categorize occupations as *replaceable* by 2012, meaning "their work could have been replaced, completely or in part, by robots" (p. 4). The authors find some evidence that the use of robots leads to reduced hours worked among low- and middle-skilled workers, with no effect for high-skilled workers. Robot densification is also found to have a positive effect on average wages and total factor productivity.

Using the IFR survey data of robot suppliers across 50 countries between 1993 and 2014, Acemoglu and Restrepo (2020) similarly explore the effect of robot adoption on labor markets. They used data from select European countries[7] to estimate global technological advances in 17 industries. They then used this data to assess how robot adoption across industries in the U.S. affects employment and wages, among other outcomes. They find that the increase in the use of robots between 1990 and 2007 is associated with a statistically significant reduction in the average employment-to-population ratio in a CZ.

Exhibit 4 on the next page, adapted from Fossen and Borgden (2022), summarizes the technology exposure measures discussed, showing differences across them.

[7] Denmark, Finland, France, Germany, Italy, Norway, Spain, Sweden, and the United Kingdom.

Exhibit 4: Key Elements of Measures of the Impact of New Technologies on Occupations.

	RTI	Computerization Probability	Task-based Computerization Probability	Advances in AI	SML	AI Exposure Score	Robot Replaceability
Source	Autor & Dorn (2009)	Frey & Osborne (2013) Automability measure	Arntz et al. (2016)	Felten et al. (2019) AIOI	Brynjolfsson et al. (2018a)	Webb (2019)	Graetz & Michaels (2018)
Time Reference	Past (1980–2005)	Next 10–20 years (viewed from 2013)	-	Past (2010–2015)	Near future (viewed from 2018)	• U.S. Census, 1960–2000 • American Community Survey (2000–2018)	14 industries in 17 countries from 1993–2007
Focus	Computerization	Computerization	Computerization	AI	ML as a subfield of AI	• Robots and software retrospectively • AI prospectively	Robots
Measurement	Merges job task requirements from Dictionary of Occupational Titles to Census occupation classifications to measure routine, abstract, and manual task content by occupation.	Experts' predictions for 71 occupations to obtain training datasets, then classifications using ML techniques.	Uses the Frey and Osborne automability measure using individual survey responses to categorize job tasks and therefore exposure to automation.	AI progress measured by the EFF mapped to 52 job requirements from O*NET and then aggregated to occupation level.	Scoring 2,069 direct work activities from O*NET through the CrowdFlower platform, then aggregated to the occupation level.	The overlap between the text of job task descriptions and the text of patents to construct a measure of the exposure of tasks to automation.	Instrumental variable uses IFR data on robot applications to classify tasks performed by robots, matched to data on 1980 U.S. occupations (prior to robots becoming ubiquitous).

Note. Adapted from "New Digital Technologies and Heterogeneous Employment and Wage Dynamics in the United States: Evidence from Individual-Level Data (Working Paper, IZA Discussion Paper No.12242)," by F. M. Fossen and A. Sorgner, 2019b, p. 41 (https://docs.iza.org/dp12242.pdf).

RESEARCH USING AUTOMATION RISK MEASURES

Several researchers have used the indicators described above in new studies. They have applied and refined measures with data from the U.S. and other countries or incorporated select measures in their analyses. In this section, we examine some of these applications and their findings.

Using Webb's (2019) methodology (e.g., methodology derived from Google Patent data and O*NET), Muro et al. (2019) show that workers with graduate or professional degrees are almost four times more exposed to AI as workers with a high school degree. Even though AI affects nearly every occupational group, white collar jobs and production workers are the groups with the most exposure among U.S. workers. Additionally, as men are more represented in occupational groups with heavy exposure (production, engineering), they are more exposed, while women's overrepresentation in occupations such as health care support and personal care correspond to a much lower exposure to AI. Similarly, the study finds that areas with economies focused on technology and products, especially denser metropolitan areas, which include both higher-tech metropolitan areas and those with a heavy manufacturing footprint, are more exposed to AI impacts.

In seeking to understand the labor effects of new technologies, Fossen and Sorgner (2019a) distinguish the effects of technological advances. They note that some will produce transformative change, changing the content of occupations without necessarily replacing workers, while others, the destructive digitalization kind, may make workers obsolete without necessarily transforming occupations. By mapping occupations across two axes (one showing Frey and Osborne's (2017) computerization probability measure, the other Felten et al.'s (2018) advances in AI), the authors identify four major occupation groups with respect to anticipated impacts. They label them "rising stars," "machine terrain," "human terrain," and "collapsing" occupations. The authors see levels of creative and social intelligence needed in an occupation as the markers of "rising stars" or "collapsing" occupations. Rising star occupations require higher levels of creative and social intelligence. In such occupations, workers are at lower risk of replacement in the near future. Such occupations may be transformed, and workers may need new qualifications to cope with the changes in the content of the occupations. Workers in collapsing occupations are more likely to be replaced, and they will need requalification.

Fossen and Sorgner (2019a) use Frey and Osborne's (2017) measure for computerization risk and Felten et al.'s (2018) AIOI measure to map occupations based on transformative/destructive dimensions. By scoring 751 occupations according to these measures and using Census data, the authors demonstrate where occupations fall in the four occupation groups—and the relative size of each in the U.S. labor market. Exhibit 5 reproduced below spatially shows how the distribution of occupations fall across the substitution/transformative dimensions. The exhibit shows large occupational clusters corresponding to the number of workers in the occupation in the collapsing quadrant, which are high on Frey and Osborne computerization probability. At the other end, rising star occupations appear at a low level of Frey and Osborne computerization probability.

Exhibit 5: Map of the Effects of Digitalization on Occupations

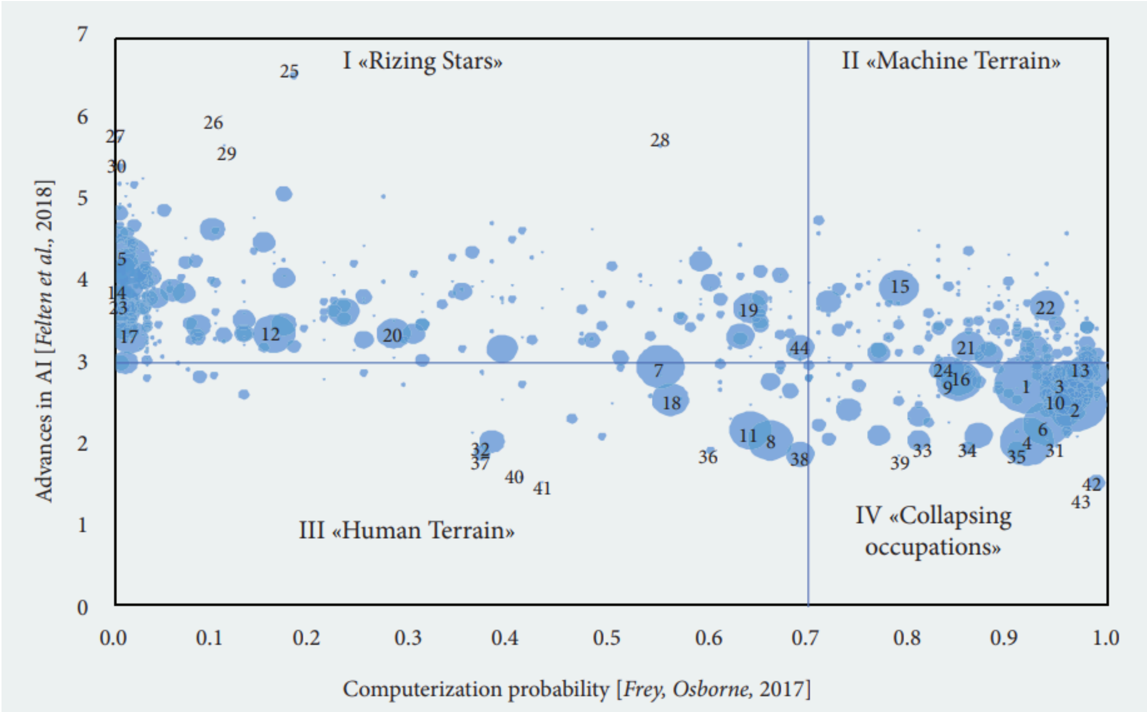

Source. From "Mapping the Future of Occupations: Transformative and Destructive Effects of New Digital Technologies on Jobs," by F. Fossen, and A. Sorgner, 2019a, *Foresight and STI Governance, 13*(2), p. 15 (https://doi.org/10.17323/2500-2597.2019.2.10.18).

Research continues to grow in this field, examining the impacts of earlier measures of exposure. Georgieff and Milanez (2021) look at what has happened to jobs deemed at high risk of automation. They use an adaptation of the Frey and Osborne (2013) model, along the lines proposed by Arntz et al. (2016), to explore the labor market outcomes in job growth and stability in these occupations as well as demographic groups most affected. They find that job growth has continued in both low- and high-risk occupations (high-risk is measured per Arntz et al. (2016) as an automatability risk of 70% or higher), but at a lower level at the latter and, in a few occupations, some decline in employment levels. On average across OECD countries, employment among the riskiest half of occupations grew by 6% compared to 18% among the least risky. Less educated workers have become increasingly concentrated in high-risk occupations. This has not been associated with a drop in the employment of low-skilled workers because their numbers have fallen relative to other education groups. However, the automation risk has grown more concentrated among the lower education group. Finally, they also find evidence that occupational-level tenure has fallen more in occupations at high risk of automation. This decline in stability is more pronounced among older workers (Georgieff & Milanez, 2021).

Studies discussed so far analyze datasets that combine survey data on the prevalence of occupations in the population with data, typically from O*NET, on the types of activities workers perform in each occupation. There are other ways to capture changes in the types of tasks and relative skill levels in jobs across the U.S. economy.

In a recent working paper, Acemoglu et al. (2022) use a unique dataset to measure impacts to job establishment (the firms listing jobs), job skill composition, wages, and employment growth. The authors built an establishment-level dataset with detailed occupation and skill information using online job listings provided by Burning Glass on the near universe of all online vacancies in the United States since 2010.

Acemoglu et al. (2022) use the AI exposure measures in the literature (AIOE, SML, and Webb's (2019) AI exposure score) as they study shifts in the types of skills and tasks described in job postings over time, examining variations across establishments. They find a strong association in the AI exposure measure and subsequent AI activity, as seen in growth in AI-related vacancies. This occurs when AIOE and AI exposure scores are used, but the association is less clear with SML. Both measures are also associated with changes in the types of skills sought in posted jobs, which the authors interpret as indicative of human-performed tasks being replaced and new tasks being generated as demonstrated by the demands for new skills. The study also finds that AI exposure is associated with lower hiring both overall and in non-AI-related vacancies. When the authors examine employment and wages at the industry- or occupation-level as opposed to the establishment-focused analyses, they do not find any relationship with employment and wages, which they explain as a result of AI adoption being too small relative to the U.S. job market to enable detection of these effects. Still, they argue their findings are suggestive of a worker displacement impact from the adoption of AI rather than a complementary impact.

Fossen and Sorgner (2019b) also examine the impacts of new technologies on individual-level employment and wages in the United States from 2011 to 2018. Using exposure measures, including Frey and Osborne's (2017) computerization probabilities of occupations, AIOI (Felten et al., 2019), and SML (Brynjolfsson & Mitchell, 2017), the study finds disparate effects on wages and employment depending on the level of occupational exposure to new technologies recorded in individual survey responses to the BLS Current Population Survey and its Annual Social and Economic Supplement[8] from January 2011 to October 2018. The analysis shows that higher risk of computerization to one's occupation (using Frey and Osborne's measure) correlates to higher probability of transition from paid employment to non-employment and to higher probability of switching to a new occupation. Occupations with higher exposure also show decreases in annual wage growth. The analysis indicates that individuals can attenuate these effects by switching to other occupations. The study also shows that exposure to AI has the opposite effect: a lower likelihood of transition to unemployment or to a new occupation, with more narrow differences between men and women. The authors maintain that these findings corroborate the hypotheses that the effect of computerization is labor substitution whereas the effect of AI is labor complementarity. Exhibit 6 illustrates these findings.

[8] The authors note that the "supplement to the survey, which is always conducted in March, contains information on various categories of income, in contrast to the interviews in the other months" (p. 12).

Exhibit 6: Estimated Effects of Digitalization on Wage Growth by Education

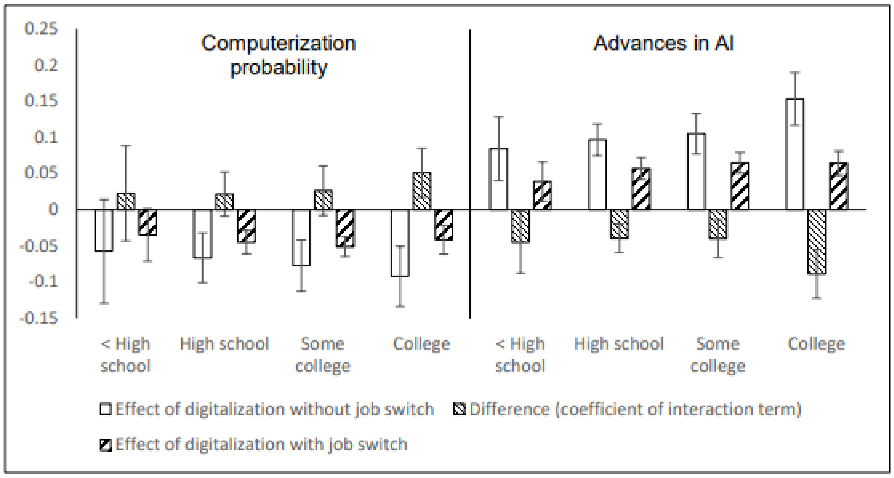

Notes from the authors: The figure illustrates the estimated coefficients from Table 6 of the effect of digitalization on wage growth, its interaction term with the dummy variable for job switchers, and the sum of these two coefficients, which indicates the total effect of digitalization for those who switch their occupations. The error bars depict 95% confidence intervals based on standard errors robust to clustering at the level of occupations. Source: Own calculations based on the annual ASEC 2011-18.

Note. From "New Digital Technologies and Heterogeneous Employment and Wage Dynamics in the United States: Evidence from Individual-Level Data (Working Paper, IZA Discussion Paper No.12242)," by F. Fossen and A. Sorgner, 2019b, p. 46 (https://docs.iza.org/dp12242.pdf).

FACTORS THAT INFLUENCE ADOPTION OF NEW TECHNOLOGIES

Up to this point, this literature scan has described research on the impact of new technology on the U.S. labor market. As shown, over the past decade researchers have developed novel methodologies to understand exposure of various occupations to advances in robotics, AI, and software applications. With access to data on robotics implementation at both the individual and establishment levels, progress has been made in measuring adoption impacts of different types of technologies.

In the research covered in this review, most researchers have been careful to note that their measures estimate potential for automation; they do not predict actual automation but the potential for automation. This is because numerous factors go into the decision to adopt new technologies as evidenced by previous waves of modernization going back to at least the industrial revolution. As an example, see the recent analysis of steam power diffusion in France in the 1800s by Ridolfi et al. (2022).

Most researchers are careful to emphasize that measures of exposure are not predictive of replacement as many issues factor into a decision to automate work, including making investments in technology and enacting shifts for the types of workers needed to operate new equipment. Indeed, the robotics and computerization literature emphasizes how the drop in costs in past decades has made possible the widespread adoption of computers within the workplace, as well as the more recent and limited industrial robot use within the United States.

Frey and Osborne (2017) note the relative costs of capital versus labor and that technology may only partially automate a job. The results may, therefore, be thought of as a measure of what is "technologically feasible" and not so much as a measure of what is economically feasible (Arntz et al., 2016, p. 10). In seeking to answer the question "Why are there still so many jobs?," Autor

(2015) notes that while automation has certainly substituted for labor, "automation also complements labor, raises output in ways that lead to higher demand for labor, and interacts with adjustments in labor supply" (p. 5). These dynamics make it difficult to predict the exact impact of new technologies, but research also makes it clear that technological advances change the types of jobs available and what those jobs pay.

For a closer-to-the ground analysis of dissemination and adoption of new technologies, Agrawal et al. (2019) use their direct experience observing industry operations at the Creative Destruction Lab at the University of Toronto to describe how "most applications of artificial intelligence have multiple forces that impact jobs, both increasing and decreasing the demand for labor. The net effect is an empirical question and will vary across applications and industries" (p. 34). In addition, they use industry cases to describe how advances in ML prediction technology may affect labor. Their cases cover a broad variety of industries—radiology, drug development, and autonomous cars. Their detailed analysis describes how even labor-substituting tasks create downstream positive effects for workers via lower costs that increase demand and efficiencies that facilitate more output.

As an example of how technological advances may lead to higher labor demand, Brynjolfsson et al. (2018b) show how higher-quality, AI-powered translation by eBay led to a higher volume of trade conducted on eBay. They show that the use of the translation tool led to an increase in U.S. exports to Spanish-speaking Latin American countries by 17.5% to 20.9%, with a potentially significant positive impact for U.S. workers.

Autor et al. (2020) similarly emphasize that there is no evidence to suggest we are moving toward a jobless future. Rather, they raise concerns regarding continued job market polarization driven by technological change in the United States. They write:

> "The causes of labor market polarization are well understood. The movement of labor from agriculture to industry to services over the 20th century has slowly eroded demand for physical labor and raised the centrality of cognitive labor in practically every walk of life. The past four decades of computerization have extended the reach of this process by displacing workers from performing routine, codifiable cognitive tasks (e.g., bookkeeping, clerical work, and repetitive production tasks) that are now readily scripted with computer software and performed by inexpensive digital machines." (p. 17)

Autor et al. (2020) are especially concerned that U.S. workers have been uniquely disadvantaged in the process when compared to other advanced economies. They show how certain mid-skill tasks will continue to be needed and how certain mid-skill occupations, especially in health care, offer middle-income salaries. They recommend targeting training investments toward such occupations (such as respiratory therapists, dental hygienists, and clinical laboratory technicians) without stopping training for those jobs in decline as some demand for these jobs will remain. They also note that we are still living with the effects and continued diffusion from technological developments of earlier decades. Newer developments in AI, ML, robotics, and additive manufacturing (3-D printing) will transform the economy, but will involve "thousands of innovations from managers, organizations, and business models" that are just barely underway.

IMPROVING FEDERAL DATA TO FACILITATE RESEARCH ON SKILL CHANGE OVER TIME

In 2017, the National Academies of Sciences, Engineering, and Medicine published a report on the activities of its ad hoc committee examining the possible impacts of automation and other IT applications on the U.S. workforce. In the report, the committee advocates for better ways to evaluate and track technology progress to help measure its impact on the workforce and to inform public policy strategies to address changes. The committee focuses on the need for research "tracking and mapping changing labor and skills demands in specific industries and occupational fields over time, along with regional variations" (p. 11).

In a report prepared for BLS, Marlar (2020) identifies major data gaps in DOL's data collection efforts that relate to researchers' ability to better understand transformations in the U.S. labor market related to new technologies. With regards to BLS's own data collection, she notes that while the Current Population Survey can generate statistics on job growth and unemployment rate by occupation, it is unable to track measures linked to job market demand flows, specifically new hires or job vacancies. These would provide clearer access to changes as they happen. Research reviewed in this literature scan shows some use of private data sources on vacancies to enable this tracking (e.g., Acemoglu et al., 2022).

Marlar (2020) also notes another gap corresponding to lack of labor market outcomes analyses by the task composition of individual jobs, such as comparing outcomes of cognitive and non-cognitive skills performed at the job. These differentiations are typically conducted by researchers to assess many phenomena of interest to policymakers and researchers, such as U.S. job market polarization (Autor et al., 2003) or individual reskilling and employment and wages growth (e.g., Georgieff & Milanez, 2021).

The third major gap identified by Marlar (2020) refers to the dearth of DOL-funded data collection focused on types of technology used and the tasks performed by individual technologies, which makes understanding the effect of technology on labor market more complicated.

Marlar (2020) provides recommendations to address each of these gaps with varying costs and benefits. First, she recommends that BLS redesign data collection to allow for time series estimates of occupational employment and compensation as one way to address the first challenge. To address the second challenge, she recommends that BLS expand JOLTS to include occupational details annually. JOLTS reports on vacancies by broad industry sectors without reference to occupations, skills, or tasks reported under each, which would require a greatly expanded sample size. To address the third challenge, she proposes a new cohort for the BLS National Longitudinal Survey of Youth (NLSY20) with a larger sample. NLSY20 contains measures of cognitive and non-cognitive skills. The larger sample would enable viable summary statistics by two-digit occupational codes. Additionally, she recommends a new module to be collected on an annual basis on tasks performed by workers to assess change as respondents age. Every 10 years, adults assessed in the original cohorts as teenagers would be reassessed for cognitive and non-cognitive skills using a validated psychometric measure of intelligence or cognitive ability.

Recognizing the centrality of tasks to occupations and to our ability to understand the impact of technologies on jobs, Marlar (2020) assesses O*NET, which as demonstrated in this literature scan is a common source of information on occupation task composition. Her findings are similar to the National Academies' ad hoc committee's (National Academies of Sciences, Engineering, and Medicine, 2017) and GAO's reports (Government Accountability Office, 2019), holding that O*NET lacks a consistent production schedule so that the occupations updates fluctuate, making it difficult to track changes in occupational skill requirements over time. She ultimately recommends that BLS develop a new task database, suggesting multiple data sources for use in creating this database. She emphasizes that, unlike O*NET's inventory of tasks, the new database should use tasks as the organizing unit. Additionally, the task database should describe tasks performed by humans and machines, allowing for an analysis of automation risk.

CONCLUSION

DOL is interested in learning more about how researchers measure skills and competencies and how they change in the workforce over time. DOL wants to know more about how to understand, identify, and anticipate these changes in connection with technological advances, leveraging new data sources and revamping old ones.

In this literature scan, we reviewed research that examines how exposed various occupations are to new and different types of technologies, with implications for employment and wage growth in different industries and jobs. We described select published research that examines the susceptibility to and impact of technological change on employment in the United States, with a focus on recent technological advances. The research findings described show that different technological advances will affect different occupations and workers. At the same time, because most jobs have a mix of tasks that can be automated and those that cannot, the potential impacts of automation are uncertain; therefore, outcomes vary and can be difficult to predict.

In very broad terms, the research shows that the latest AI progress will affect highly skilled positions, whereas previous waves of computerization and robotics dissemination affected low- and middle-skilled positions. Jobs that involve interpersonal tasks and creativity are less likely to be automated as these tasks are not well-suited to replacement by AI.

To understand and anticipate these effects, researchers have leveraged multiple data sources to establish the prevalence of tasks and skills within occupations or jobs, the susceptibility of these tasks to replacement by technology, the growth of these occupations and wages within these occupations, and variations in demographic groups in these trends. Exhibit 7 summarizes data sources used by researchers discussed in this scan, detailing the information each data source provides about occupational task/skill composition and technology applications to tasks and occupations.

Exhibit 7: Summary of Data Sources in the Field

Dimensions	Data Sources	Detail
Occupational tasks and skills composition	O*NET and its predecessor, the Dictionary of Occupational Titles	O*NET provides information regarding work characteristics, experience requirements, job responsibilities, and the state of the labor market
	Burning Glass vacancy listings	Near universe of online vacancies in the U.S. since 2010 (e.g., Acemoglu et al., 2022); AI skills measured in Burning Glass 2010 data (e.g., Felten et al., 2019)

Dimensions	Data Sources	Detail
	NLSY20	Includes measures of cognitive and non-cognitive skills, but sample does not enable valid statistics by occupation (Marlar, 2020)
	OECD Programme for the International Assessment of Adult Competencies (PIAAC)	Measures adults' proficiency in key information-processing skills; also asks questions about how often an individual's job usually involves particular tasks (e.g., Arntz et al., 2016)
Technological applications	Google Patents Public Data	Includes title, abstract, and Cooperative Patent Classification codes (e.g., Webb, 2019)
	Expert coding	ML researchers from the Oxford University Engineering Sciences Department subjectively hand-labelled 70 occupations, then expanded dataset using ML (e.g., Frey & Osborne, 2017)
	IFR industry data	Compiles information from national robot federations on industrial robots, including deliveries, prices (e.g., Acemoglu & Restrepo, 2020; Graetz & Michaels, 2018)
	EFF AI Progress Measurement dataset	Tracks reported progress on metrics of AI performance across separate AI applications, such as image recognition, speech recognition, translation, or abstract strategy games, drawing on data from multiple sources, including academic literature, review articles, blog posts, and websites focused on AI (Felten et al., 2019)

As noted in the previous section, these data sources are not always well-suited to capturing skill demand changes over time. To explore a role for DOL in developing or collaborating to develop data sources to advance this research field, in the fall 2022 CEO and ETA will host a research roundtable with recognized experts to learn more about state-of-the-art data sources and methodologies in the field and to gather their recommendations.[9]

[9] Four researchers participated in the virtual roundtable, which took place on October 21, 2022: Dr. Robert Seamans from New York University, Dr. Ina Ganguli from the University of Massachusetts Amherst, Dr. Ben Zweig from Revelio Labs, and Dr. Daniel Rock from the University of Pennsylvania.

REFERENCES

Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, *40*(S1). https://doi.org/10.1086/718327

Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from U.S. labor markets. *Journal of Political Economy*, *128*(6), 2188–2244. https://doi.org/10.1086/705716

Agrawal, A., Gans, J. S., & Goldfarb, A. (2019). Artificial intelligence: The ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives, 33*(2), 31–50. DOI: 10.1257/jep.33.2.31

Arntz, M., Gregory, T., & Zierahn, U. (2016). The risk of automation for jobs in OECD countries. *OECD Social, Employment and Migration Working Papers*. https://doi.org/10.1787/5jlz9h56dvq7-en

Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, *29*(3), 3–30. https://doi.org/10.1257/jep.29.3.3

Autor, D. H., & Dorn, D. (2009). *This job is 'getting old:' Measuring changes in job opportunities using occupational age structure* (Working Paper 14652). National Bureau of Economic Research. https://www.nber.org/system/files/working_papers/w14652/w14652.pdf

Autor, D. H., & Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the U.S. labor market. *American Economic Review*, *103*(5), 1553–1597. https://doi.org/10.1257/aer.103.5.1553

Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, *118*(4), 1279–1333. https://doi.org/10.1162/003355303322552801

Autor, D. H., Mindell, D., & Reynolds, E. (2020, November 17). *The work of the future: Building better jobs in an age of Intelligent Machines*. MIT Work of the Future. Retrieved August 1, 2022, from https://workofthefuture.mit.edu/research-post/the-work-of-the-future-building-better-jobs-in-an-age-of-intelligent-machines/

Brynjolfsson, E., Hui, X., & Liu, M. (2018b). Does machine translation affect international trade? Evidence from a large digital platform. *National Bureau of Economic Research Working Papers, 24917*. https://doi.org/10.3386/w24917

Brynjolfsson, E., & Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science*, *358*(6370), 1530–1534. https://doi.org/10.1126/science.aap8062

Brynjolfsson, E., Mitchell, T., & Rock, D. (2018a). What can machines learn and what does it mean for occupations and the economy? *AEA Papers and Proceedings*, *108*, 43–47. https://doi.org/10.1257/pandp.20181019

Felten E.W., Raj M., Seamans R. (2018) A Method to Link Advances in Artificial Intelligence to Occupational Abilities. American Economic Association Papers and Proceedings, vol. 108, pp. 54–57.

Felten, E. W., Raj, M., & Seamans, R. (2019). The occupational impact of artificial intelligence: Labor, skills, and polarization. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3368605

Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, *42*(12), 2195–2217. https://doi.org/10.1002/smj.3286

Fossen, F., & Sorgner, A. (2019a). Mapping the future of occupations: Transformative and destructive effects of new digital technologies on jobs. *Foresight and STI Governance*, *13*(2), 10–18. https://doi.org/10.17323/2500-2597.2019.2.10.18

Fossen, F. M., & Sorgner, A. (2019b). *New digital technologies and heterogeneous employment and wage dynamics in the United States: Evidence from individual-level data* (Working Paper, IZA Discussion Paper No.12242). https://docs.iza.org/dp12242.pdf

Frey, C. B., & Osborne, M. (2013). *The future of employment*. Oxford Martin Programme on Technology and Employment. https://ora.ox.ac.uk/objects/uuid:4ed9f1bd-27e9-4e30-997e-5fc8405b0491/download_file?safe_filename=future-of-employment.pdf&file_format=application%2Fpdf&type_of_work=journal+article

Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, *114*, 254–280. https://doi.org/10.1016/j.techfore.2016.08.019

Georgieff, A., & Milanez, A. (2021). What happened to jobs at high risk of automation? *OECD Social, Employment and Migration Working Papers*, *255*. https://doi.org/10.1787/10bc97f4-en

Government Accountability Office. (2019). *Workforce automation: Better data needed to assess and plan for effects of advanced technologies on jobs*. (GAO Publication No. 19-257). https://www.gao.gov/products/gao-19-257#summary_recommend

Graetz, G., & Michaels, G. (2018). Robots at work. *The Review of Economics and Statistics*, *100*(5), 753–768. https://doi.org/10.1162/rest_a_00754

Mann, K., & Püttmann, L. (2021). Benign effects of automation: New evidence from patent texts. *The Review of Economics and Statistics*, 1–45. https://doi.org/10.1162/rest_a_01083

Marlar, J. (2020). *Assessing the impact of new technologies on the labor market: Key constructs, gaps, and data collection strategies for the Bureau of Labor Statistics*. U.S. Bureau of Labor Statistics. https://www.bls.gov/bls/congressional-reports/assessing-the-impact-of-new-technologies-on-the-labor-market.htm

Muro, M., Whiton, J., & Maxim, R. (2019, November 20). *What jobs are affected by AI? Better-paid, better-educated workers face the most exposure*. Brookings Institute. https://www.brookings.edu/research/what-jobs-are-affected-by-ai-better-paid-better-educated-workers-face-the-most-exposure/

National Academies of Sciences, Engineering, and Medicine. (2017). *Information technology and the U.S. workforce: Where are we and where do we go from here?* The National Academies Press. https://doi.org/10.17226/24649

Ridolfi, L., Salvo, C., & Weisdorf, J. (2022). *The impact of mechanisation on wages and employment: Evidence from the diffusion of steam power*. CEPR Discussion Paper No. DP17389. https://ssrn.com/abstract=4144721

U.S. Department of Labor. (n.d.). *About us*. Retrieved April 16, 2023, from https://www.dol.gov/general/aboutdol

Webb, M. (2019). The impact of artificial intelligence on the labor market. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3482150

Webb, M. (2020). *The impact of artificial intelligence on the labor market*. https://www.michaelwebb.co/webb_ai.pdf