

RESEARCH ARTICLE

Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses

Edward Felten¹ | Manav Raj²  | Robert Seamans² 

¹Princeton University, Princeton, New Jersey

²NYU Stern School of Business, New York, New York

Correspondence

Robert Seamans, NYU Stern School of Business, 44 West 4th Street, New York, New York, 10012, USA.

Email: rseamans@stern.nyu.edu

Abstract

Research Summary: We create and validate a new measure of an occupation's exposure to AI that we call the AI Occupational Exposure (AIOE). We use the AIOE to construct a measure of AI exposure at the industry level, which we call the AI Industry Exposure (AIIE) and a measure of AI exposure at the county level, which we call the AI Geographic Exposure (AIGE). We also describe several ways in which the AIOE can be used to create firm level measures of AI exposure. We validate the measures and describe how they can be used in different applications by management, organization and strategy scholars.

Managerial Summary: Although artificial intelligence (AI) promises to spur economic growth, there is widespread concern that it could displace workers, alter industry trajectories, and reshape organizations. Despite the interest in this area, we have limited ability to study the effects of AI on occupations, firms, industries, and geographies because of limited availability of data that measures exposure to AI. To address this limitation, we create and validate a new measure of an occupation's exposure to AI that we call the AI Occupational Exposure (AIOE). We use the AIOE to construct a measure of AI exposure at the industry level (AIIE)

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2021 The Authors. *Strategic Management Journal* published by John Wiley & Sons Ltd.

and county level (AIGE). We describe how our measures can be useful to scholars and policy-makers interested in identifying the effect of AI on markets.

KEYWORDS

artificial intelligence, industry, innovation, occupation, technology

1 | INTRODUCTION

Recent advances in artificial intelligence (AI) have generated excitement about AI's potential to spur economic growth, and scholars believe that AI has the potential to be “the most important general-purpose technology of our era” (Brynjolfsson & McAfee, 2017). However, there is concern that advances in AI may also have significant consequences for labor markets, firms, and industries by displacing workers, transforming occupational jurisdictions, altering strategy, and affecting performance. For decades now, scholars have considered whether and how rapid advances in information technologies change the nature of competition and strategy (Bennett & Hall, 2020; Bettis & Hitt, 1995; Tippins & Sohi, 2003). In recent years, researchers have increasingly begun to investigate how AI affects firm design, strategy, organizational learning, and management (e.g., Balasubramanian, Xu, & Ye, 2020; Bughin, Kretschmer, & van Zeebroeck, 2019; Iansiti & Lakhani, 2020; Jia, Luo, & Fang, 2020a, 2020b; Khashabi & Kretschmer, 2019; Raj & Seamans, 2019; Wuebker, Saouma, & McGahan, 2018). However, despite great interest in the academic literature and public press about AI's effects on occupations, firms, and markets, there has been little systematic collection of evidence. Part of the reason for the lack of evidence is that the rapid advancement in AI is a nascent phenomenon, and accordingly, appropriate tools to measure its impact have yet to be developed (McElheran, 2018; Raj & Seamans, 2018).

In order to fill this gap, we develop a new measure of the exposure to AI across occupations that we call the AI Occupational Exposure (AIOE).¹ This measure links common and general applications of AI to workplace abilities and occupations. We show the potential of the occupation-level AIOE by using it to construct two derivative measures. We aggregate the AIOE to the industry level to construct an AI Industry Exposure (AIIE) as well as to the county-level to construct a geographic measure of AI exposure that we call the AI Geographic Exposure (AIGE). We also describe several ways in which the AIOE can be used to create firm level measures of AI exposure. We validate our AIOE measure and then describe how these measures can be used in different applications by strategy, management, and innovation scholars. We have made these datasets freely available for use by scholars, policy-makers, and practitioners.²

We build our measures by linking common AI applications to occupational abilities using a crowd-sourced dataset. We aggregate the effect at the ability level to construct a measure that identifies the potential exposure of occupations to AI. While similar in spirit to recent work (e.g., Brynjolfsson, Mitchell, & Rock, 2018; Frey & Osborne, 2017; Georgieff & Milanez, 2021; Tolan et al., 2020; Webb, 2020), our methodology is unique in that it links specific applications of AI to occupational abilities and is agnostic as to whether AI substitutes for or

¹The AIOE measure is a modified version of a method described by Felten, Raj, and Seamans (2018).

²The datasets can be accessed via GitHub: <https://github.com/AIOE-Data/AIOE>.

complements workplace abilities and occupations. We discuss differences with existing datasets in more detail below. We believe our methodology has the potential to be used in a variety of applications within the field of strategy and, as we describe below, can be used to construct AI exposure at the firm level as well (see, e.g., Acemoglu, Autor, Hazell, & Restrepo, 2020).

Our article contributes in several ways. First, we develop a new methodology linking applications of AI technology to human abilities. We use this method to generate occupation-level, industry-level, and geographic-level “AI Exposure” measures, which are available for other researchers to use. By describing the data construction and validation in the body of this article and providing this data for use by other researchers, we are answering a call from Ethiraj, Gambardella, and Helfat (2019) for more articles that develop and describe datasets for other researchers to use. Second, we situate our measures within existing data by comparing the AIOE with other measures that link work and occupations to AI, automation, and robotics (e.g., Brynjolfsson, Frank, Mitchell, Rahwan, & Rock, 2020; Frey & Osborne, 2017; Webb, 2020). Third, we describe several ways in which the measures can be used by researchers to study the effects of exposure to AI on workers, firms, industries, and geographies.

The article proceeds as follows. In the next section, we outline the methodology that we use to construct the AI Occupational Exposure (AIOE) and its accompanying AI Industry Exposure (AIIE) and AI Geographic Exposure (AIGE) measures. In the third section, we describe the steps we take to validate our methodology and resulting measures. In the fourth section, we describe potential applications for the datasets, with particular attention to the use of the AIOE for organizational level studies. The fifth section concludes.

2 | CONSTRUCTION OF THE AI EXPOSURE MEASURES

Artificial intelligence is a construct with varying definitions and interpretations: “general” AI refers to computer software that can think and act on its own, which does not yet exist; “narrow” AI refers to computer software that relies on highly sophisticated algorithmic techniques to find patterns in data and make predictions about the future (Broussard, 2018). Because narrow AI algorithms “learn” from existing data to improve performance, these techniques are often referred to as “machine learning.” To construct our measure of occupational exposure to AI, we start by identifying different applications of these machine learning techniques and then link these applications to occupational abilities.

Our methodology isolates the exposure to specific functions of AI and creates a measure of the aggregate exposure to AI at the occupation level by looking at the abilities used within an occupation. We consider “labor” as the bundle of skills and abilities that are used within a specific occupation (D. Autor & Handel, 2013; Cohen, 2012), and this “micro” approach allows us to examine how AI may be related to the abilities that each occupation requires. We link specific functions of AI to different types of abilities in the US workforce using a crowdsourced dataset; then, considering the work content of each occupation, we aggregate the ability-level exposure to AI applications at the occupation level. Note that our measure of exposure does not attempt to measure whether AI is a complement to or a substitute for labor, but rather how likely it is that the occupation is exposed to AI in some way.

Our analytical strategy borrows in part from Felten et al.’s (2018) methodology, which links progress in AI applications to construct a backward-looking measure that evaluates the effect of AI on occupations from 2010 to 2015. We focus on a similar set of AI applications to build our measure, but instead of accounting for historical progress to construct a backward-looking

measure, we take a forward-looking approach to identify occupational exposure to AI. We also modify Felten et al.'s (2018) methodology to scale the aggregate exposure to AI at the occupation level by the portfolio of abilities used within an occupation. This modification allows us to account for occupation scope and more accurately measure the relative occupation-level exposure to AI.

2.1 | Data

To construct our measure, we identify a set of common and well-developed applications of AI based on categories defined and described by the Electronic Frontier Foundation (EFF) AI Progress Measurement project and connect these applications of AI to data on occupational abilities from the Occupational Information Network (O*NET) database developed by the United States Department of Labor.

The Electronic Frontier Foundation (EFF) is a well-regarded digital rights nonprofit that was founded in 1990, focusing on issues related to digital rights and privacy. Searches in Google Scholar show thousands of results, including articles written by the EFF as well as articles that cite EFF's work in a variety of fields, including law, economics, technology, and others (e.g., Casadesus-Masanell & Hervas-Drane, 2015; Eaton, Elaluf-Calderwood, Sorensen, & Yoo, 2015; Farrell & Shapiro, 2008; Liebowitz, 2006; Lu, Wang, & Bendle, 2019).

As one of its activities, the EFF collects and maintains statistics about the progress of AI across separate artificial intelligence applications. For each application, the EFF monitors progress in the field by drawing on data from multiple sources, including academic literature, review articles, blog posts, and websites focused on artificial intelligence. For all information collected, the EFF only includes data from verified sources that document proof of their findings (AI Progress Measurement, 2019), and the EFF AI metrics have been cited in various computer science outlets (e.g., Carlson, 2019; Croeser & Eckersley, 2019; Krinitskiy et al., 2018). For the purposes of this study, we focus on the 10 AI applications for which the EFF has recorded measured scientific activity and progress in the technology from 2010 onward as we believe these are the applications experiencing the fastest growth and most likely to be used in the medium-term.³ While we use the EFF categorization to identify applications of AI, we note that EFF's categorization is consistent with the categorization of AI applications in other outlets (e.g., Bessen, Impink, Reichensperger, & Seamans, 2018; Martínez-Plumed et al., 2020; Perrault et al., 2019).

The O*NET data define and describe professions in the modern-day American workplace (O*NET Database Releases Archive at O*NET Resource Center, 2019) and are frequently used to measure occupational work or task content in academic research (e.g., Autor & Handel, 2013; Brynjolfsson et al., 2018; Goos, Manning, & Salomons, 2009). For the purposes of our analysis, we focus on the 52 distinct abilities O*NET uses to describe the workplace activities of each occupation. O*NET notes both the importance and the level or prevalence of each ability within an occupation (using a 1–5 and 1–7 scale, respectively).

³There has been less activity in the EFF AI Progress Measurement project since 2017. However, we believe that the 10 applications we have selected are appropriate representations of common applications of AI technology. We discuss how the EFF categorizations compare to other sources and explore the robustness of the measure to using a subset of applications in Online Appendix C.

TABLE 1 EFF application definitions

AI application	Definition
Abstract strategy games	The ability to play abstract games involving sometimes complex strategy and reasoning ability, such as chess, go, or checkers, at a high level.
Real-time video games	The ability to play a variety of real-time video games of increasing complexity at a high level.
Image recognition	The determination of what objects are present in a still image.
Visual question answering	The recognition of events, relationships, and context from a still image.
Image generation	The creation of complex images.
Reading comprehension	The ability to answer simple reasoning questions based on an understanding of text.
Language modeling	The ability to model, predict, or mimic human language.
Translation	The translation of words or text from one language into another.
Speech recognition	The recognition of spoken language into text.
Instrumental track recognition	The recognition of instrumental musical tracks.

Note: The considered applications are those that have experienced meaningful scientific progress as measured by the Electronic Frontier Foundation at the time of the draft. Definitions are based on those used by the Electronic Frontier Foundation.

2.2 | Methodology

2.2.1 | Selection of AI applications

We construct the AIOE using data from 10 selected AI applications: abstract strategy games, real-time video games, image recognition, visual question answering, image generation, reading comprehension, language modeling, translation, speech recognition, and instrumental track recognition. This set of applications does not comprehensively cover the set of applications for which AI could ultimately be used; however, based on our conversations with field experts, we believe that these represent fundamental applications of AI that are likely to have implications for the workforce and are applications that cover the most likely and most common uses of AI.⁴ Table 1 presents the definition of each application provided by the EFF.

2.2.2 | Linking AI applications to O*NET abilities

To link the AI applications to workplace abilities, for each of the AI applications considered in our analysis (abstract strategy games, real-time video games, image recognition, visual question answering, image generation, reading comprehension, language modeling, translation, and speech recognition), we use a crowd-sourced data set constructed using survey responses of “gig workers” from Amazon’s Mechanical Turk (mTurk) web service. Through this survey, we generate a measure of application-ability relatedness for each combination bound between 0 and 1 (survey methodology and measure construction is discussed in more detail in Appendix A). We organize this measure of application-ability relatedness into a matrix that connects the 10 EFF AI

⁴Robustness checks around the selection of AI applications are presented in Online Appendix C.

applications to the 52 O*NET occupational abilities. Using the measures of application-ability relatedness from this matrix, we then calculate an ability-level exposure as follows:

$$A_{ij} = \sum_{i=1}^{10} x_{ij} \quad (1)$$

In this equation, i indexes the AI application and j indexes the occupational ability. The ability-level exposure, A , is calculated as the sum of the 10 application-ability relatedness scores, x , as constructed using mTurk survey data. For example, if the application-ability relatedness score between an occupational ability and each separate AI application were 0.5, the ability-level exposure would be 5.0. By calculating the ability-level AI exposure as a sum of all the AI applications, we are equally weighting each application. Some may argue with this choice, especially given that certain applications appear related (e.g., visual question answering and image recognition). Although we choose this approach for simplicity and to avoid making arbitrary decisions about whether and how to weight applications, we show in Appendix C that our measures are largely consistent if we construct our AIOE measure using alternative sets of applications, suggesting that weighting the applications is unlikely to have a meaningful impact on the measure.⁵ Our approach also assumes that each application has an independent effect on an ability and does not consider interactions across applications. Although this is a limitation, considering the myriad ways in which the applications could interact would be infeasible.

2.2.3 | Calculating occupational exposure

We next use the O*NET occupational definitions to evaluate the exposure to the AI technology for each occupation. We rely on the O*NET 24.3 database released in May 2020, as it is the most up-to-date version of the O*NET data at the time of analysis. Our occupation-level AIOE measure is constructed as follows:

$$AIOE_k = \frac{\sum_{j=1}^{52} A_{ij} \times L_{jk} \times I_{jk}}{\sum_{j=1}^{52} L_{jk} \times I_{jk}} \quad (2)$$

In this equation, i indexes the AI application, j indexes the occupational ability, and k indexes the occupation. A_{ij} represents the ability-level exposure score calculated in Equation 1. We weight the ability-level AI exposure by the ability's prevalence (L_{jk}) and importance (I_{jk}) within each occupation as measured by O*NET by multiplying the ability-level AI exposure by the prevalence and importance scores for that ability within each occupation, scaled so that they are equally weighted.⁶ These prevalence and importance scores allow us to properly account for

⁵This is because we find that, across the different applications of AI, the relationship between AI and abilities is strongest for “cognitive” abilities.

⁶To equally weight prevalence and importance, we divide each score by the maximum value O*NET considers for the metric.

the presence of different abilities within an occupation. Abilities that are integral to an occupation have high prevalence and importance scores, while those that are used less often or are less vital have lower prevalence and importance scores. For example, for CEOs, O*NET considers abilities such as oral comprehension (importance: 4.5; prevalence: 4.88) and oral expression (importance: 4.38; prevalence: 5) as highly important, and abilities such as speed of limb movement and static strength as nonessential (for both, importance: 1; prevalence: 0).⁷

We measure an occupation's aggregate exposure to AI by summing this weighted ability-level AI exposure across all abilities in an occupation. In an adjustment to Felten et al.'s (2018) methodology, we scale the aggregated exposure to AI across all abilities by the weighted sum of the prevalence and importance of all abilities used in the occupation to account for the total required ability set within an occupation. This scaling provides us with a measure of the *relative exposure* to AI. We believe this adjustment is important owing to the nature of the O*NET occupational definitions. Some occupations as defined by O*NET have many abilities that are considered important and prevalent, while others have a smaller number. Because an occupation's *aggregate exposure* to AI is measured by summing across abilities, occupations that use more abilities are likelier to have higher levels of aggregate exposure. Without the adjustment based on the weighted sum of all abilities, our methodology would overweight the exposure to AI for broader occupations that require more abilities.

To clarify this with an example, consider two separate occupations—surgeons and physicists. Both occupations score highly on *aggregate exposure* to AI (calculated by summing the ability-level AI exposure across all abilities). Using the outlined approach, surgeons have the third highest aggregate AI exposure in our sample (3.28 SDs above the mean), while physicists have the fourth highest aggregate AI exposure in our sample (3.14 SDs above the mean). However, although the aggregate exposure of AI may be similar for these two occupations, the *relative exposure* to AI within the occupation differs greatly because each of these occupations rely on a different portfolio of abilities. In particular, surgeons rely on a broad range of abilities from all ability families — cognitive, physical, psychomotor, and sensory abilities — while physicists rely heavily on a smaller set of largely cognitive abilities. To account for such differences across occupations, we scale the aggregate exposure of AI across an occupation by the breadth of abilities required in that occupation. After that adjustment, we find that physicists are still considered highly exposed to AI (the 61st ranked occupation with an AIOE measure 1.35 SDs above the mean), while surgeons are no longer classified as highly exposed (the 387th ranked occupation with an AIOE measure of 0.05 SDs below the mean). The adjustment, therefore, accounts for the different scope of occupations and more accurately measures the relative exposure to AI within an occupation.

The O*NET data are organized according to each occupation's eight-digit SOC classification; we collapse the AIOE to occupations at the six-digit SOC level to align with other major data sources. Because the vast majority of eight-digit SOC classifications contain only one six-digit SOC classification, this change does not alter our sample size greatly. We collapse the 832 occupations at the eight-digit SOC level to 774 occupations at the six-digit SOC level by calculating the mean AIOE across all occupations defined by O*NET within the six-digit SOC level. Finally, we standardize our AIOE measures such that the mean across occupations is zero and the SD is one.⁸ Ultimately, our methodology produces a score that measures how the most prevalent applications of AI are related to occupations at the six-digit SOC level based on their ability

⁷There is a high correlation between an occupation's prevalence and importance scores. The AIOE measure is largely consistent using either prevalence or importance alone to weight the ability-level AI exposure.

⁸This step results in occupations with negative AIOE measures. We do not intend to suggest that occupations can have negative exposure to AI; rather this transformation just allows for ease of comparison across occupations.

composition as defined by O*NET. We again emphasize that this does not constitute a measure for how substitutable or automatable an occupation is, as we are agnostic about whether or when AI will augment or replace human labor.

2.2.4 | AIIE and AIGE

We construct our industry-level measure of AI exposure by aggregating the AIOE across occupations within an industry. We construct an AI Industry Exposure (AIIE) by taking a weighted average of the AIOE using industry employment based on the four-digit NAICS classification in 2019, the latest data available at the time of writing. This measure allows us to track which industries have the most exposure to AI. We construct our geographic-level measure of AI exposure by aggregating the AIIE across industries within a geographic area using county-level employment in 2019 (we use the US county designations [FIPS codes] to construct measures at the county and state levels). In addition to providing us with a measure of AI exposure at industry and geographic levels, this exercise demonstrates the flexibility of our methodology. Our measures can be updated based on the changing occupational descriptions and definitions provided by the O*NET database, changing occupational composition within an industry, or changing occupational composition within a geographic area. In addition, we believe that this measure can be aggregated to other levels, such as the firm or sub-industry level, using the distribution of employment or wages across occupations within a grouping.

We provide the datasets created by our methodology in the Appendices available via GitHub.⁹ A list of occupations and the associated AIOE measures is in Appendix A; a list of industries at the four-digit NAICS classification measuring industry exposure to AI technology is in Appendix B; and the AIGE scores, measuring geographic exposure to AI technology, are in Appendix C. We include the matrix connecting AI applications to occupational abilities in Appendix D and the standardized ability-level AI exposure in Appendix E. In the following section, we analyze and validate the AIOE measures.

3 | VALIDATION OF THE AI OCCUPATION EXPOSURE SCORES

Owing to the novel nature of our methodology and the limited research in this area, validation of the metric presents challenges since there are no clear benchmarks to compare with the AIOE measures. It is unclear how one might expect the exposure to AI to manifest itself in occupations while remaining agnostic about its effect on occupations and patterns of adoption, and most existing metrics for measuring occupation-level exposure to AI seek to determine the likelihood of substitution (Brynjolfsson et al., 2018; Frey & Osborne, 2017).

Given these challenges, we validate the AIOE measures in several ways. First, as an initial “gut check,” we qualitatively discuss the differences across occupations that receive the highest and lowest scores in our sample. We show that exposure to AI seems to be highest in white-collar occupations and industries. Second, we provide a deeper discussion for several occupations, including some that are often considered to be threatened by advances in AI such as taxi drivers and long-haul truck drivers, to provide a sense of how advances in AI are related to specific occupational abilities. Third, we present robustness checks for the construction of our

⁹The datasets can be accessed via GitHub: <https://github.com/AIOE-Data/AIOE>.

matrix that connects AI applications to abilities and show that our scores are mostly consistent even if we exclude different applications of AI. Fourth, we use data from job postings to present evidence that the AIOE is positively associated with the use of AI skills within an occupation.¹⁰

3.1 | Most- and least-exposed occupations and industries

To begin to understand the nature of the AIOE measure, we first compare the highest-scoring occupations with the lowest-scoring occupations. These represent the occupations with the most and least reliance on abilities that are likely to be affected by advances in AI technology. Table 2 presents the 20 occupations with the highest and lowest AIOE measures in our sample.

The highest-scoring occupations (i.e., the ones that the AIOE predicts are most exposed to advances in AI technologies) consist almost entirely of white-collar occupations that require advanced degrees, such as genetic counselors, financial examiners, and actuaries. The lowest-scoring occupations are largely nonoffice jobs that require a high degree of physical effort and exertion, such as dancers, fitness trainers and aerobics instructors, and painters and plasterers.

We can similarly compare the highest- and lowest-scoring industries based on our AIIE measure to identify which industries our methodology suggests have the most exposure to AI technologies. Table 3 presents the 20 industries with the highest- and lowest-scoring AIIE measure in our sample.

Just as with the ranked occupations, we find that the most exposed industries tend to be white-collar industries requiring high levels of education, such as financial services, accounting, insurance, and legal services. Perhaps not surprisingly, exposure to AI is particularly high in certain service industries that involve a high level of information processing. On the other hand, the lowest-scoring industries tend to be blue-collar industries that involve manual labor, such as support activities for crop production, building and dwellings services, construction contracting services, and warehousing and storage. Intuitively, these occupation and industry rankings reflect (1) our aim to isolate the exposure to advances in AI (as opposed to, say, robotics, machine vision, autonomous guided vehicles, or other types of advanced technologies), and (2) the abilities most likely to be affected by AI. Whereas robotic technologies often involve physical manipulation and are capable of performing complex manual tasks, AI technologies are largely software-based and rely on iterative learning and perception (Raj & Seamans, 2019). Accordingly, we expect AI to have a limited influence on the role of physical abilities in occupations and industries. Instead, our methodology leads us to believe that AI is likely to have the biggest impact on abilities related to information processing. The AIOE and AIIE measures reflect the larger influence of AI on cognitive tasks and abilities and a smaller influence on physical tasks and abilities. We believe that this relationship is in line with current scientific perceptions of AI technologies (e.g., Brynjolfsson & McAfee, 2014; Brynjolfsson & Mitchell, 2017) and accordingly provides a measure of support for our methodology.

3.2 | Discussion of selected occupations

We next turn to a discussion of selected occupations. We start by comparing and contrasting the scores of two occupations that require some similar abilities but have disparate AI exposure

¹⁰While we present the highest- and lowest-scoring industries based on AIIE, we largely focus on validating our AIOE measure rather than the AIIE measure since the AIIE is constructed using the AIOE.

TABLE 2 Occupations with the highest and lowest AIOE measures

Rank	Highest scoring	Lowest scoring
1	Genetic counselors	Dancers
2	Financial examiners	Fitness trainers and aerobics instructors
3	Actuaries	Helpers—painters, paperhangers, plasterers, and stucco masons
4	Purchasing agents, except wholesale, retail, and farm products	Reinforcing iron and rebar workers
5	Budget analysts	Pressers, textile, garment, and related materials
6	Judges, magistrate judges, and magistrates	Helpers—Brickmasons, Blockmasons, stonemasons, and tile and marble setters
7	Procurement clerks	Dining room and cafeteria attendants and bartender helpers
8	Accountants and auditors	Fence erectors
9	Mathematicians	Helpers—roofers
10	Judicial law clerks	Slaughterers and meat packers
11	Education administrators, postsecondary	Landscaping and Groundskeeping workers
12	Clinical, counseling, and school psychologists	Athletes and sports competitors
13	Financial managers	Fallers
14	Compensation, benefits, and job analysis specialists	Structural iron and steel workers
15	Credit authorizers, checkers, and clerks	Cement masons and concrete finishers
16	History teachers, postsecondary	Terrazzo workers and finishers
17	Geographers	Rock splitters, quarry
18	Epidemiologists	Plasterers and stucco masons
19	Management analysts	Brickmasons and Blockmasons
20	Arbitrators, mediators, and conciliators	Roofers

Note: Occupations are ranked by their constructed AIOE measure at the six-digit Standard Occupational Classification (SOC) level. See draft for a detailed description of the construction of the AIOE measure. 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.

scores—surgeons and meat slaughterers—to highlight the importance of cognitive abilities in determining occupational exposure to AI. We then compare two occupations that are similar in the presence of cognitive abilities but are have dissimilar AOIE scores—mathematical technicians and accountants and auditors—to examine what other facets of an occupation may affect exposure to AI. In Appendix B, we also review two occupations frequently discussed in the context of AI—truck drivers and taxi drivers.

3.2.1 | Surgeons and slaughterers

We attempt to provide more context for the AIOE measures by comparing two professions that share some similarities abilities required yet have highly disparate AIOE measures—surgeons and meat slaughterers. Superficially, we might expect these occupations to have similar AIOE measures, as both require deft physical manipulation of human or animal tissue. Although the occupations require similar physical abilities, such as manual dexterity, finger dexterity,

TABLE 3 Industries with the highest and lowest AIIE measures

Rank	Highest scoring	Lowest scoring
1	Securities, commodity contracts, and other financial investments and related activities	Support activities for crop production
2	Accounting, tax preparation, bookkeeping, and payroll services	Services to buildings and dwellings
3	Insurance and employee benefit funds	Foundation, structure, and building exterior contractors
4	Legal services	Animal slaughtering and processing
5	Agencies, brokerages, and other insurance related activities	Building finishing contractors
6	Nondepository credit intermediation	Warehousing and storage
7	Other investment pools and funds	Fiber, yarn, and thread Mills
8	Insurance carriers	Support activities for rail transportation
9	Software publishers	Sawmills and wood preservation
10	Lessors of nonfinancial intangible assets (except copyrighted works)	Support activities for water transportation
11	Agents and managers for artists, athletes, entertainers, and other public figures	Logging
12	Credit intermediation and related activities (5,221 and 5,223 only)	Other specialty trade contractors
13	Computer systems design and related services	Waste collection
14	Management, scientific, and technical consulting services	Postal service (federal government)
15	Monetary authorities-central Bank	Highway, street, and bridge construction
16	Office administrative services	Truck transportation
17	Other information services	Apparel knitting Mills
18	Data processing, hosting, and related services	Seafood product preparation and packaging
19	Business schools and computer and management training	Local messengers and local delivery
20	Grantmaking and giving services	Utility system construction

Note: Industries are ranked by their constructed AIIE Measure at the four-digit North American Industry Classification System (NAICS) level. See draft for a detailed description of the construction of the AIIE Measure. Industry titles are taken from the Bureau of Labor Statistics. 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.

and arm-hand steadiness, the occupations' AIOE measures suggest that surgeons are far more exposed to AI than slaughterers. The AIOE measure for surgeons is at the 52nd percentile in relation to other occupations in our sample, while the AIOE measure for slaughterers is at the 2nd percentile (indeed, it is the 10th least-exposed occupation).

The difference between the AIOE measures in these two occupations seems to arise from the cognitive abilities required by each occupation. Although the two occupations require similar physical abilities, a number of cognitive abilities related to problem solving, such as problem sensitivity, deductive and inductive reasoning, and information ordering, are highly important for surgeons but not for meat slaughterers. Our methodology suggests that occupations'

exposure to AI largely stems from these cognitive abilities; accordingly, surgeons have a much higher AIOE measure than meat slaughterers.

Generalizing from this specific example, our AIOE measure emphasizes that the presence of cognitive abilities plays a large role in determining how exposed an occupation is to AI. This has implications for how work content and processes may be affected by advances in AI. Literature has argued that computer-based technologies may have a large effect on cognitive tasks, substituting for routine cognitive tasks and complementing nonroutine cognitive tasks (Autor, Levy, & Murnane, 2003). The AIOE measure suggests that, like previous iterations of computer-based technologies, AI is likely to disproportionately affect cognitive tasks and occupations. We expect occupations that require a greater amount of problem solving, logical reasoning, and perception to be more exposed to AI than occupations that largely require physical abilities.

3.2.2 | Mathematical technicians and accountants and auditors

As discussed above, we find that the primary driver of the AIOE is the presence of cognitive abilities within an occupation. However, because our methodology relies on granular abilities, we can look beyond just the presence of cognitive abilities to examine what other occupational characteristics are related to exposure to AI. To do this, we consider two occupations that require a similar level of cognitive abilities but have dissimilar AIOE scores—mathematical technicians and accountants and auditors.

Based on O*NET's occupational definitions and categorization of abilities, 80.0% of abilities required for both these occupations (weighted by importance and level) are cognitive abilities. These are both occupations that could be considered relatively “cognitive”, as for the median occupation, 56.8% of abilities required are cognitive abilities. However, despite a similar level of cognitive abilities in these two occupations, there is a disparity in the AIOE scores. Accountants and auditors have an AIOE score at the 99th percentile relative to other occupations in our sample, while mathematical technicians have an AIOE score at the 68th percentile.

The difference is accounted for by the presence of sensory versus physical or psychomotor abilities. Cognitive abilities are more exposed to AI relative to all other abilities; however, differences remain across noncognitive abilities. In particular, we find that sensory abilities, defined as abilities that influence visual, auditory, and speech perception, are more exposed to AI technologies than physical or psychomotor abilities. This is perhaps not surprising given that AI technologies are particularly known to be well-suited for tasks involving classification, categorization, and pattern recognition (Broussard, 2018; Choudhury, Starr, & Agarwal, 2020; Kotsiantis, Zaharakis, & Pintelas, 2006), and perception may be keenly involved in such tasks. Of the noncognitive abilities required for accountants and auditors, 90.6% are sensory abilities, while for mathematical technicians, only 52.6% are sensory abilities. The higher relative importance of sensory abilities (vs. physical or psychomotor abilities) for accountants and auditors results in the higher exposure to AI relative to mathematical technicians, despite a similar level of cognitive abilities.

3.3 | Quantitative validation

In addition to our qualitative review of occupations and their AIOE measures, we test the robustness of our AIOE measure and seek to validate the measure quantitatively. We describe these tests briefly here and discuss them in more detail in Appendix C. We test the robustness

of our measure along four separate dimensions: (1) we consider the construction of our matrix linking AI applications to occupational abilities and show that the matrix connecting AI applications and O*NET abilities remains consistent using a subset of respondents with graduate degrees or with computer science or engineering degrees; (2) we test the sensitivity of the measure to which set of AI applications is included and show that our measure is robust to alternative application selection, as all AI applications are most related to cognitive and sensory abilities; (3) we use job postings to show a strong positive correlation between the AIOE measure and the use of AI skills within an occupation; and (4) we examine how changes in occupational definitions by O*NET affect the AIOE measure and show that the measure is relatively consistent over short- to medium-term time horizons.

4 | POTENTIAL APPLICATIONS OF AIOE, AIIE, AND AIGE

We believe our AIOE, AIIE, and AIGE measures can be used by a range of scholars interested in occupational, industry, firm, and regional outcomes. Thus, our measure should be useful to scholars investigating research questions in multiple areas including competitive and corporate strategy, organizational design, human resource management, industrial organization, geography of innovation, regional and business dynamism, and entrepreneurship. In this section, we describe how the AIOE, AIIE, and AIGE measures can be used by scholars to better understand how exposure to AI is affecting occupations, firms, industries, and regions. We first describe distinctions between the measures and existing datasets at a high-level to inform of particularly suitable uses of the data before providing some examples of research questions that the data could be used to examine.

4.1 | Distinctions between the AIOE and existing datasets

The AIOE and related measure are distinct from existing datasets in a number of ways. First, our methodology considers specific applications of AI (e.g., image recognition, speech recognition, and others) and links them to workplace abilities and then to occupations, industries, and geographies, rather than considering the effect of automation (Frey & Osborne, 2017; Mann & Püttmann, 2017) or AI more broadly (Brynjolfsson et al., 2018, 2020; Webb, 2020) on labor, without specifying whether automation occurs via AI, robots, sensors, or another type of technology. Second, our measure identifies the relative *exposure* to AI but remains agnostic as to whether AI complements or substitutes for tasks and labor (e.g., Brynjolfsson et al., 2020; Frey & Osborne, 2017; Webb, 2020). Third, our approach provides a snapshot of occupational exposure to AI based on the current nature of occupations, rather than relying on expert projections or crowdsourced evaluations to identify which tasks may be suitable to AI or automation or how the task composition of an occupation may change as it is affected by AI (Brynjolfsson et al., 2018, 2020; Frey & Osborne, 2017).

Such distinctions inform us of particularly suitable uses of the data. Given that our measure considers AI (and applications of AI) specifically, we believe our measure is best suited for studies that focus on AI exclusively rather than automation or robotics more broadly. Because our measure seeks to measure exposure to AI rather than identify abilities or occupations that are likely to be replaced or automated, we also believe that our measure can be useful for scholars who aim to explore the trade-off between substitution and complementarities, as it does not seek to identify abilities or occupations that are likely to be replaced or automated. Further, because the AIOE, AIIE, and AIGE are forward-looking measures, we believe they can inform

us which occupations, industries, and geographies are most likely to be exposed to AI based on current occupational definitions.

4.2 | Potential uses of the AIOE, AIIE, and AIGE across fields

4.2.1 | Competitive and corporate strategy

The AIOE can also be used to construct a measure of AI exposure at the firm-level. To do this, a researcher would need the occupational breakdown of employees within firms. The researcher could use the occupation names or codes to link our AIOE measures to the occupations in the firm, and then created an AI exposure score for the firm by weighting the AIOE measure for each occupation present in the firm by the percentage of the firm's employees in that occupation. Several existing datasets allow for this kind of firm-level data construction. For example, the Occupational Employment Survey (OES), which is produced by the BLS, provides publicly available occupational data at the industry level. The underlying micro-data, at the firm level, are confidential but available to BLS-approved researchers, and the OES micro-data have been used to study occupational differences across firms by Ayyagari and Maksimovic (2017), Handwerker (2020) and Ma, Ouimet, and Simintzi (2019). As another example, data from the Equal Employment Opportunity Commission (EEOC) allow researchers to study occupational-level details within firms (e.g., Ferguson & Koning, 2018; Koning & Ferguson, 2019; Tomaskovic-Devey et al., 2016). Finally, the Burning Glass Technologies (BG) data that we use in one of our validation exercises can be used to study occupational differences across firms. Indeed, Acemoglu et al. (2020) match our AIOE measure to firm level BG data to study how a firm's AI exposure affects its hiring of AI skilled and non-AI skilled workers. Interestingly, Acemoglu et al. also do the same for the Brynjolfsson et al. (2018) and Webb (2020) measures and discuss the differences across the three measures.

A firm-level measure of AI exposure could be used to answer a number of questions related to competitive strategy. Strategy and management scholars have often sought to identify the different firm characteristics that predict novel technology adoption or diffusion (e.g., Brynjolfsson & McElheran, 2016; Greve, 2009; Majumdar & Venkataraman, 1998; Souder, Zaheer, Sapienza, & Ranucci, 2017). We believe that a firm-level construction of our AI measure could be used to identify what characteristics predict firm-level adoption of AI technologies given similar exposure. In a recent article, Adner, Puranam, and Zhu (2019) call for additional research on digital innovation and its implications for firm strategy. A burgeoning line of research has already begun to study how AI affects internal organization and strategy (e.g., Bughin et al., 2019; Iansiti & Lakhani, 2020; Jia et al., 2020a; Khashabi & Kretschmer, 2019). We believe that our measures could be used by such scholars to identify firms that are facing greater exposure to novel AI technologies and study changes in organization or strategy.

4.2.2 | Human capital

We believe that our AIOE measure can be a valuable tool for scholars seeking to study how advances in AI relate to individual- and occupation-level outcomes within firms. Articles in the public press frequently discuss the threat that developments in AI may pose to workers

(e.g., Jordan, 2018; Stark, 2017; Wacker, 2017), and a stream of research has begun to examine how AI and automation may affect the workforce (e.g., Acemoglu & Restrepo, 2018; Arntz, Gregory, & Zierahn, 2016; Frey & Osborne, 2017; Mann & Püttmann, 2017). However, despite such interest in the press and in the academic community, there is very little systematic evidence on the impact of AI on labor. Because the AIOE measure is constructed to be agnostic to the effect of AI on labor demand, we believe that the AIOE can be a useful tool for understanding the circumstances in which AI substitutes for or complements labor.

Beyond this fundamental question regarding labor substitution, the AIOE can be used to study the relationship between AI, human capital, and labor productivity. For example, previous literature has suggested that equipment automation does not substitute for the value of human capital, at least within the semiconductor industry (Hatch & Dyer, 2004), but it is unclear whether and how the emergence of AI will alter the value of human capital within firms. While capital-labor substitution has traditionally been driven by highly specialized technologies (Istvan, 1992), AI is a general-purpose technology (Goldfarb, Taska, & Teodoridis, 2019) with a wide range of specialized applications that could challenge this common belief. Our measures may be useful to scholars seeking to extend research seeking to understand how AI and other digital or automating technologies may affect human productivity (Adler, Benner, Brunner, Macduffie, & Osono, 2009; Brynjolfsson & McElheran, 2016; Choudhury et al., 2020; Cowgill, 2019; Furman & Teodoridis, 2020). Finally, there are many other human-capital related issues of interest to management scholars, including career mobility (Karim & Williams, 2012), team building and team learning (Huckman, Staats, & Upton, 2008; Shah, Agarwal, & Echambadi, 2019), the use of virtual teams (Maznevski & Chudoba, 2000), and others. The AIOE can be used to study how AI affects organizations and managers along all these dimensions.

4.2.3 | Organizational design

The rapid emergence of AI is likely to alter the distribution of value created by employees across an organization (Brynjolfsson & McAfee, 2017; Bughin et al., 2019; Fountaine, McCarthy, & Saleh, 2019; Iansiti & Lakhani, 2020). Previous research has examined how ICTs have affected firm boundaries and transaction costs (McElheran & Forman, 2019; Vakili & Kaplan, 2019). Similarly, scholars could use the AIOE to study the links between AI and occupational characteristics and work content, to better understand the implications of AI for organizational design and structure. For example, Raj and Seamans (2019) call for more studies on how AI changes the nature of work and the resulting implications for organizational design, structure, and boundaries.

4.2.4 | Industry evolution

Technological change can have meaningful consequences for industry structure and lead to the entry or exit of firms within an industry (e.g., Abernathy & Utterback, 1978; Agarwal & Gort, 2002; Klepper, 1996, 1997). Historically, product innovation has led to industry shakeouts and changes in firm innovative strategy (Klepper & Simons, 2000; Utterback, 1987), and already evidence has emerged that AI has affected talent recruitment, healthcare, and legal industries (Cowgill, 2019; Galasso & Luo, 2018; Ramesh, Kambhampati, Monson, & Drew, 2004). Literature has suggested that the increased prevalence of ICTs may decrease competition and increase


concentration within industries (D. Autor, Dorn, Katz, Patterson, & Van Reenen, 2017; Bessen, 2017; Van Reenen, 2018). Recent work has drawn links between high-technology sectors and declining business dynamism and has suggested that the rise of ICTs may be partially responsible for declining productivity growth (Bijnens & Konings, 2018; Haltiwanger, Hathaway, & Miranda, 2014). Other work suggests, however, that an increase in exposure to software availability increases industry entry and also increases exit by the oldest and most established firms in the industry (Bennett & Hall, 2020). As AI technology continues to advance and diffuse, understanding how the AI affects industry growth and evolution becomes increasingly important. For example, Babina, Fedyk, He, and Hodson (2020) study how investments in AI affect industry dynamics and find that increased adoption of AI technologies is associated with an increase in industry concentration. We believe the AIIE measure could be used by researchers to study the links between AI and industry concentration, firm entry, or business dynamism.

4.2.5 | Agglomeration and location-based advantages

Finally, we believe the AIGE measure can be used to study the relationship between AI and regional level outcomes. Previous research has found that technology spillovers tend to be geographically localized because transfer and diffusion of (often tacit) knowledge benefits from local interactions, whether it be through communication, collaboration, or localized human capital (Alcacer & Chung, 2014; Almeida & Kogut, 1999; Audretsch & Feldman, Audretsch & Feldman, 1996; Boschma, 2005; Shaver & Flyer, 2000). The extent to which local interactions occur and whether they will eventually result in meaningful advances in technology may depend upon regional capabilities that govern the innovation processes. These capabilities could include infrastructure, natural resources, institutions, universities, firms, skills, and culture (Cooke, 2001; Maskell & Malmberg, 1999). For example, local infrastructure such as roads and railways can enhance regional innovation by increasing diffusion of technological knowledge among local inventors (Agrawal, Galasso, & Oettl, 2017; Perlman, 2015). It is possible that a region's exposure to AI generates tacit knowledge about how to work with and benefit from the technology that in turn constitutes a capability upon which future innovation will build. Moreover, one might expect that the effect is greater in areas where there are more interactions, perhaps because of better roads and infrastructure or more foot traffic (Roche, 2019). On the other hand, to the extent that AI substitutes for labor, on net, then a region's exposure to AI will indicate future job losses and economic decline.

As a simple example of the potential of the AIGE, in Figure 1, we present a map of the continental United States that displays the AIGE by county. More highly exposed counties (i.e., those with relatively more employment in occupations with high AIOE measures) are shown in red and less exposed counties in blue.

Figure 1 presents some immediate patterns for AI exposure across the US. It appears that, on average, urban counties are more highly exposed to AI than rural counties. For example, the Atlantic corridor, comprising Boston, New York City, Philadelphia, Baltimore, and Washington, D.C., has a relatively high-level of AI exposure, as does the Bay Area in California. While rural areas have lower AI exposure on average, there are some exceptions. For example, the state of Iowa has a number of counties with relatively high exposure to AI despite being predominantly rural. We present this figure as a simple example of the potential of the data and



Note: The charted color represents the county-level AIGE measure winsorized at the 1st and 99th percentile.

FIGURE 1 AIGE in the continental United States.

Note: The charted color represents the county-level AIGE measure winsorized at the 1st and 99th percentile

leave it to future scholars to disentangle the underlying causes and effects of such heterogeneity in exposure.

5 | CONCLUSION

Despite excitement about AI's benefits to organizations and concern about its potential effect on labor, there is little systematic evidence of either. Our study makes several contributions toward understanding these effects. First, we describe and validate a new methodology for linking advances in AI to human abilities. Second, using this methodology, we develop a measure of an occupation's exposure to AI, which we call the AI Occupational Exposure (AIOE); a measure of an industry's exposure to AI, which we call the AI Industry Exposure (AIIE); and a measure of a geographic area's exposure to AI, which we call the AI Geographic Exposure (AIGE). We also discuss ways in which our AIOE measure can be aggregated to the firm level. We expect the new measures will be useful to researchers interested in the link between AI and occupation, firm, industry, and region-oriented outcomes. Third, we validate the constructed measures and discuss their potential uses.

We believe that these datasets can be useful to policymakers and scholars seeking to understand how AI will affect individuals, firms, and markets. The methodology used to construct the data is dynamic and can be updated as O*NET updates occupational definitions over time or occupational composition changes across industries and geographies. The AIOE, AIIE, and AIGE can potentially be used by policymakers to identify the occupations, industries, and geographic areas that are most likely to be affected by future advances. These measures can be used by scholars in economics,

management, and strategy to answer questions regarding the effect of AI on firm, market, and local outcomes.

Our article is closely related to a small number of others that provide their own measures of the exposure to new technologies on occupations and industries (Brynjolfsson et al., 2018, 2020; Frey & Osborne, 2017; Tolan et al., 2020; Webb, 2020). We believe that our measures are unique in that they focus on specific applications of AI technologies rather than considering AI or automation broadly; they are agnostic as to whether the effect of AI serves as a complement or a substitute; and they provide a snapshot of occupational exposure to AI based on the current nature of the occupation. We believe our approach complements existing approaches and can be useful to researchers, depending on the specific questions being asked. Some researchers have already used the AIOE, including Fossen and Sorgner (2019), who study how advances in AI create opportunities for growth-oriented entrepreneurship, Goldfarb et al. (2019), who use the AIOE to examine whether AI has the characteristics of a general-purpose technology within the healthcare industry, and Acemoglu et al. (2020) who use the AIOE to study how a firm's AI exposure affects hiring patterns. It is our hope that future work can continue to investigate how the exposure to AI manifests across occupations, geographies, and backgrounds.

AI technologies are rapidly advancing and are expected to have dramatic effects on the economy. Given the sizable effect that we suspect AI will have on firms and markets, it is critical to have measures that allow us to study this important phenomenon. While much work must still be done to examine the relationship between AI and labor, we believe our measures can be useful tools in this endeavor and that our study contributes to the growing body of literature examining this important topic.

ACKNOWLEDGEMENTS

The authors are grateful to participants of the 2018 American Economic Association Annual meeting, 2019 Temple University Conference on AI, ML, and Business Analytics, 2019 Munich Summer Institute, 2019 EMAEE conference and the 2021 AI & Strategy Consortium, and seminar participants at Copenhagen Business School, Cornell Dyson School of Applied Economics and Management, Georgia Tech Sheller School of Business, Loughborough University London, NYU Stern School of Business, UC Boulder Leeds School of Business, UCL School of Management, and the UMass-Lowell Manning School of Business. We also thank Dan Sands and Derek Choe for thoughtful comments and feedback. All errors are our own.

OPEN RESEARCH BADGES



This article has been awarded Open Data Badge for making publicly available the digitally-shareable data necessary to reproduce the reported results. Data is available at <https://github.com/AIOE-Data/AIOE>.

ORCID

Manav Raj  <https://orcid.org/0000-0002-3657-7843>

Robert Seamans  <https://orcid.org/0000-0001-9058-4190>

REFERENCES

- Abernathy, W. J., & Utterback, J. M. (1978). Patterns of industrial innovation. *Technology Review*, 80(7), 40–47.
- Acemoglu, D., & Restrepo, P. (2018). Robots and Jobs: Evidence from US Labor Markets. Working Paper, 90.

- Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2020). AI and Jobs: Evidence from Online Vacancies. MIT Working Paper.
- Adler, P. S., Benner, M., Brunner, D. J., Macduffie, J. P., & Osono, E. (2009). Perspectives on the productivity dilemma. *Journal of Operations Management*, 27(2), 99–113.
- Adner, R., Puranam, P., & Zhu, F. (2019). What is different about digital strategy? From quantitative to qualitative change. *Strategy Science*, 4(4), 253–261. <https://doi.org/10.1287/stsc.2019.0099>
- Agarwal, R., & Gort, M. (2002). Firm and product life cycles and firm survival. *American Economic Review*, 92(2), 184–190. <https://doi.org/10.1257/000282802320189221>
- Agrawal, A., Galasso, A., & Oettl, A. (2017). Roads and innovation. *The Review of Economics and Statistics*, 99(3), 417–434.
- Alcácer, J., & Chung, W. (2014). Location strategies for agglomeration economies. *Strategic Management Journal*, 35(12), 1749–1761. <https://doi.org/10.1002/smj.2186>
- Almeida, P., & Kogut, B. (1999). Localization of knowledge and the mobility of engineers in regional networks. *Management Science*, 45(7), 905–917.
- Arntz, M., Gregory, T., & Zierahn, U. (2016). The risk of automation for jobs in OECD countries. <https://doi.org/10.1787/5jlz9h56dvq7-en>
- Audretsch, D. B., & Feldman, M. P. (1996). R&D spillovers and the geography of innovation and production. *The American Economic Review*, 86(3), 630–640.
- Autor, D., & Handel, M. J. (2013). Putting tasks to the test: Human capital, job tasks, and wages. *Journal of Labor Economics*, 31(S1), S59–S96. <https://doi.org/10.1086/669332>
- Autor, D., Dorn, D., Katz, L. F., Patterson, C., & van Reenen, J. (2017). Concentrating on the fall of the labor share. *American Economic Review*, 107(5), 180–185. <https://doi.org/10.1257/aer.p20171102>
- 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.
- Ayyagari, M., & Maksimovic, V. (2017). Fewer and less skilled? Human capital, competition, and entrepreneurial success in manufacturing (SSRN Scholarly Paper ID 3093443). *Social Science Research Network*. <https://doi.org/10.2139/ssrn.3093443>
- Babina, T., Fedyk, A., He, A. X., & Hodson, J. (2020). Artificial intelligence, firm growth, and industry concentration (SSRN Scholarly Paper ID 3651052). *Social Science Research Network*. <https://doi.org/10.2139/ssrn.3651052>
- Balasubramanian, N., Xu, M., & Ye, Y. (2020). Faster but more fragile? Organizational learning in the presence of machine Learning. Working Paper.
- Bennett, V. M., & Hall, T. A. (2020). Software availability and entry. *Strategic Management Journal*, 41(5), 950–962. <https://doi.org/10.1002/smj.3105>
- Bessen, J. E. (2017). Industry Concentration and Information Technology (SSRN scholarly paper ID 3044730). *Social Science Research Network*. <https://papers.ssrn.com/abstract=3044730>
- Bessen, J. E., Impink, S. M., Reichensperger, L., & Seamans, R. (2018). The business of AI startups (SSRN Scholarly Paper ID 3293275). *Social Science Research Network*. <https://doi.org/10.2139/ssrn.3293275>
- Bettis, R. A., & Hitt, M. A. (1995). The new competitive landscape. *Strategic Management Journal*, 16(S1), 7–19. <https://doi.org/10.1002/smj.4250160915>
- Bijnens, G., & Konings, J. (2018). Declining business dynamism and information technology. VoxEU.Org. <https://voxeu.org/article/declining-business-dynamism-and-information-technology>
- Boschma, R. (2005). Proximity and innovation: A critical assessment. *Regional Studies*, 39(1), 61–74. <https://doi.org/10.1080/0034340052000320887>
- Broussard, M. (2018). *Artificial unintelligence: How computers misunderstand the world*. Cambridge, MA: MIT Press.
- Brynjolfsson, E., & McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. New York, NY: W. W. Norton & Company.
- Brynjolfsson, E., & McAfee, A. (2017). The business of artificial intelligence. *Harvard Business Review*. <https://hbr.org/2017/07/the-business-of-artificial-intelligence>
- Brynjolfsson, E., & McElheran, K. (2016). The rapid adoption of data-driven decision-making. *American Economic Review*, 106(5), 133–139. <https://doi.org/10.1257/aer.p20161016>
- 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. (2018). 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>
- Brynjolfsson, E., Frank, M. R., Mitchell, T., Rahwan, I., & Rock, D. (2020). Quantifying the Impact of Machine Learning on Work. Working Paper.
- Bughin, J. R., Kretschmer, T., & van Zeebroeck, N. (2019). Experimentation, learning and stress: The role of digital Technologies in Strategy Change (SSRN Scholarly Paper ID 3328421). *Social Science Research Network*. <https://doi.org/10.2139/ssrn.3328421>
- Carlson, K. W. (2019). Safe artificial general intelligence via distributed ledger technology. *Big Data and Cognitive Computing*, 3(3), 40. <https://doi.org/10.3390/bdcc3030040>
- Casadesus-Masanell, R., & Hervas-Drane, A. (2015). Competing with privacy. *Management Science*, 61(1), 229–246. <https://doi.org/10.1287/mnsc.2014.2023>
- Choudhury, P., Starr, E., & Agarwal, R. (2020). Machine learning and human capital complementarities: Experimental evidence on bias mitigation. *Strategic Management Journal*, 41(8), 1381–1411. <https://doi.org/10.1002/smj.3152>
- Cohen, L. E. (2012). Assembling jobs: A model of how tasks are bundled into and across jobs. *Organization Science*, 24(2), 432–454. <https://doi.org/10.1287/orsc.1110.0737>
- Cooke, P. (2001). Regional innovation systems, clusters, and the knowledge economy. *Industrial and Corporate Change*, 10(4), 945–974. <https://doi.org/10.1093/icc/10.4.945>
- Cowgill, B. (2019). Bias and Productivity in Humans and Machines. Working Paper, 30.
- Croeser, S., & Eckersley, P. (2019). Theories of parenting and their application to artificial intelligence. Paper presented at: Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society, 423–428. <https://doi.org/10.1145/3306618.3314231>
- Eaton, B., Elaluf-Calderwood, S., Sorensen, C., & Yoo, Y. (2015). Distributed tuning of boundary resources: The case of Apple's iOS service system. *Management Information Systems Quarterly*, 39(1), 217–243.
- Ethiraj, S. K., Gambardella, A., & Helfat, C. E. (2019). Articles on datasets. *Strategic Management Journal*, 40(5), 713–714. <https://doi.org/10.1002/smj.3000>
- Farrell, J., & Shapiro, C. (2008). How strong are weak patents? *American Economic Review*, 98(4), 1347–1369. <https://doi.org/10.1257/aer.98.4.1347>
- Felten, E. W., Raj, M., & Seamans, R. (2018). A method to link advances in artificial intelligence to occupational abilities. *AEA Papers and Proceedings*, 108, 54–57. <https://doi.org/10.1257/pandp.20181021>
- Ferguson, J.-P., & Koning, R. (2018). Firm turnover and the return of racial establishment segregation. *American Sociological Review*, 83(3), 445–474. <https://doi.org/10.1177/0003122418767438>
- Fossen, F. M., & Sorgner, A. (2019). The effects of digitalization of work on entry into entrepreneurship. *Academy of Management Proceedings*, 2019(1), 11095. <https://doi.org/10.5465/AMBPP.2019.11095abstract>
- Fountaine, T., McCarthy, B., & Saleh, T. (2019). Building the AI-powered organization. *Harvard Business Review*. <https://hbr.org/2019/07/building-the-ai-powered-organization>
- 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>
- Furman, J. L., & Teodoridis, F. (2020). Automation, research technology, and Researchers' trajectories: Evidence from computer science and electrical engineering. *Organization Science*, 31(2), 330–354. <https://doi.org/10.1287/orsc.2019.1308>
- Galasso, A., & Luo, H. (2018). *Punishing robots: Issues in the economics of tort liability and innovation in artificial intelligence* (pp. 493–504). An Agenda: The Economics of Artificial Intelligence.
- Georgieff, A., & Milanez, A. (2021). What happened to jobs at high risk of automation? In *OECD social, employment and migration working papers*. (No. 255; *OECD Social, Employment and Migration Working Papers*). OECD Publishing. <https://ideas.repec.org/p/oec/elsaab/255-en.html>
- Goldfarb, A., Taska, B., & Teodoridis, F. (2019). *Could machine learning be a general-purpose technology? Evidence from online job postings* (SSRN Scholarly Paper ID 3468822). Social Science Research Network. <https://doi.org/10.2139/ssrn.3468822>
- Goos, M., Manning, A., & Salomons, A. (2009). Job polarization in Europe. *American Economic Review*, 99(2), 58–63. <https://doi.org/10.1257/aer.99.2.58>
- Greve, H. R. (2009). Bigger and safer: The diffusion of competitive advantage. *Strategic Management Journal*, 30(1), 1–23. <https://doi.org/10.1002/smj.721>

- Haltiwanger, J., Hathaway, I., & Miranda, J. (2014). Declining business dynamism in the U.S. high-technology sector. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.2397310>
- Handwerker, E. W. (2020). Outsourcing, Occupationally Homogeneous Employers, and Growing Wage Inequality in the United States. BLS Working Paper. <https://www.bls.gov/osmr/research-papers/2020/ec200030.htm>
- Hatch, N. W., & Dyer, J. H. (2004). Human capital and learning as a source of sustainable competitive advantage. *Strategic Management Journal*, 25(12), 1155–1178. <https://doi.org/10.1002/smj.421>
- Huckman, R. S., Staats, B. R., & Upton, D. M. (2008). Team familiarity, role experience, and performance: Evidence from Indian software services. *Management Science*, 55(1), 85–100. <https://doi.org/10.1287/mnsc.1080.0921>
- Iansiti, M., & Lakhani, K. R. (2020). Competing in the age of AI. *Harvard Business Review*. <https://hbr.org/2020/01/competing-in-the-age-of-ai>
- Istvan, R. L. (1992). A new productivity paradigm for competitive advantage. *Strategic Management Journal*, 13(7), 525–537. <https://doi.org/10.1002/smj.4250130705>
- Jia, N., Luo, X., & Fang, Z. (2020a). Can artificial intelligence (AI) substitute or complement managers? Divergent outcomes for transformational and transactional managers in a field experiment. Working Paper.
- Jia, N., Luo, X., & Fang, Z. (2020b). Management by AI-Bot: A field experiment of employee performance coaching. Working Paper.
- Jordan, C. (2018). *Working alongside AI in tomorrow's labor market?*. The Hill. <https://thehill.com/blogs/congress-blog/technology/367900-working-alongside-ai-in-tomorrows-labor-market>
- Karim, S., & Williams, C. (2012). Structural knowledge: How executive experience with structural composition affects intrafirm mobility and unit reconfiguration. *Strategic Management Journal*, 33(6), 681–709. <https://doi.org/10.1002/smj.1967>
- Khashabi, P., & Kretschmer, T. (2019). *Digital transformation and organization design: A complex relationship* (SSRN Scholarly Paper ID 3437334). Social Science Research Network. <https://doi.org/10.2139/ssrn.3437334>
- Klepper, S. (1996). Entry, exit, growth, and innovation over the product life cycle. *American Economic Review*, 86(3), 562–583.
- Klepper, S. (1997). Industry life cycles. *Industrial and Corporate Change*, 6(1), 145–182. <https://doi.org/10.1093/icc/6.1.145>
- Klepper, S., & Simons, K. L. (2000). Dominance by birthright: Entry of prior radio producers and competitive ramifications in the U.S. television receiver industry. *Strategic Management Journal*, 21(10–11), 997–1016. [https://doi.org/10.1002/1097-0266\(200010/11\)21:10/11<997::AID-SMJ134>3.0.CO;2-0](https://doi.org/10.1002/1097-0266(200010/11)21:10/11<997::AID-SMJ134>3.0.CO;2-0)
- Koning, R., & Ferguson, J.-P. (2019). *Does public ownership and accountability increase diversity?: Evidence from IPOs* (SSRN Scholarly Paper ID 3316178). Social Science Research Network. <https://doi.org/10.2139/ssrn.3316178>
- Kotsiantis, S. B., Zaharakis, I. D., & Pintelas, P. E. (2006). Machine learning: A review of classification and combining techniques. *The Artificial Intelligence Review*, 26(3), 159–190. <https://doi.org/10.1007/s10462-007-9052-3>
- Krinitzkiy, M., Verezhenskaya, P., Grashchenkov, K., Tilinina, N., Gulev, S., & Lazzara, M. (2018). Deep convolutional neural networks capabilities for binary classification of polar mesocyclones in satellite mosaics. *Atmosphere*, 9, 426. <https://doi.org/10.3390/atmos9110426>
- Liebowitz, S. (2006). File sharing: Creative destruction or just plain destruction? *Journal of Law and Economics*, 49(1), 1–28. <https://chicagounbound.uchicago.edu/jle/vol49/iss1/1>
- Lu, S., Wang, X., & Bendle, N. (2019). Does piracy create online word of mouth? An empirical analysis in the movie industry. *Management Science*, 66(5), 2140–2162. <https://doi.org/10.1287/mnsc.2019.3298>
- Ma, W., Ouimet, P., & Simintzi, E. (2019). *Mergers and acquisitions, technological change and inequality* (SSRN Scholarly Paper ID 2793887). Social Science Research Network. <https://doi.org/10.2139/ssrn.2793887>
- Majumdar, S. K., & Venkataraman, S. (1998). Network effects and the adoption of new technology: Evidence from the U.S. telecommunications industry. *Strategic Management Journal*, 19(11), 1045–1062. [https://doi.org/10.1002/\(SICI\)1097-0266\(199811\)19:11<1045::AID-SMJ990>3.0.CO;2-0](https://doi.org/10.1002/(SICI)1097-0266(199811)19:11<1045::AID-SMJ990>3.0.CO;2-0)
- Mann, K., & Püttmann, L. (2017). Benign effects of automation: New evidence from patent texts. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.2959584>
- Martínez-Plumed, F., Tolan, S., Pesole, A., Hernández-Orallo, J., Fernández-Macías, E., & Gómez, E. (2020). Does AI qualify for the job?: A bidirectional model mapping labour and AI intensities. Paper presented at: Proceedings of the AAI/ACM Conference on AI, Ethics, and Society, 94–100. <https://doi.org/10.1145/3375627.3375831>

- Maskell, P., & Malmberg, A. (1999). Localised learning and industrial competitiveness. *Cambridge Journal of Economics*, 23(2), 167–185. <https://doi.org/10.1093/cje/23.2.167>
- Maznevski, M. L., & Chudoba, K. M. (2000). Bridging space over time: Global virtual team dynamics and effectiveness. *Organization Science*, 11(5), 473–492. <https://doi.org/10.1287/orsc.11.5.473.15200>
- McElheran, K. (2018). Economic Measurement of AI. Working Paper, 29.
- McElheran, K., & Forman, C. (2019). *Firm Organization in the Digital age: IT use and vertical transactions in U.S. manufacturing (SSRN Scholarly Paper ID 3396116)*. Social Science Research Network. <https://doi.org/10.2139/ssrn.3396116>
- O*NET Database Releases Archive at O*NET Resource Center. (2019). https://www.onetcenter.org/db_releases.html
- Perlman, E. R. (2015). Dense enough to be brilliant: Patents, urbanization, and transportation in nineteenth century America. In *CEH discussion papers (Working Paper)*. Boston University. <https://ideas.repec.org/p/auu/hpaper/036.html>
- Perrault, C. R., Shoham, Y., Brynjolfsson, E., Clark, J., Etchemendy, J., Grosz, B., ... Niebles, J. C. (2019). *The AI index 2019 annual report*. AI Index Steering Committee, Human-Centered AI Institute, Stanford University.
- AI Progress Measurement. (2019). Electronic Frontier Foundation. <https://www.eff.org/ai/metrics>
- Raj, M., & Seamans, R. (2018). *AI, labor, productivity and the need for firm-level data* (No. w24239). National Bureau of Economic Research. <https://doi.org/10.3386/w24239>
- Raj, M., & Seamans, R. (2019). Primer on artificial intelligence and robotics. *Journal of Organization Design*, 8(1), 1–14.
- Ramesh, A. N., Kambhampati, C., Monson, J. R. T., & Drew, P. J. (2004). Artificial intelligence in medicine. *Annals of the Royal College of Surgeons of England*, 86(5), 334–338. <https://doi.org/10.1308/147870804290>
- Roche, M. P. (2019). Taking innovation to the streets: Microgeography, physical structure and innovation. *The Review of Economics and Statistics*, 105(5), 1–47. https://doi.org/10.1162/rest_a_00866
- Shah, S. K., Agarwal, R., & Echambadi, R. (2019). Jewels in the crown: Exploring the motivations and team building processes of employee entrepreneurs. *Strategic Management Journal*, 40(9), 1417–1452. <https://doi.org/10.1002/smj.3027>
- Shaver, J. M., & Flyer, F. (2000). Agglomeration economies, firm heterogeneity, and foreign direct investment in the United States. *Strategic Management Journal*, 21(12), 1175–1193. [https://doi.org/10.1002/1097-0266\(200012\)21:12<1175::AID-SMJ139>3.0.CO;2-Q](https://doi.org/10.1002/1097-0266(200012)21:12<1175::AID-SMJ139>3.0.CO;2-Q)
- Souder, D., Zaheer, A., Sapienza, H., & Ranucci, R. (2017). How family influence, socioemotional wealth, and competitive conditions shape new technology adoption. *Strategic Management Journal*, 38(9), 1774–1790. <https://doi.org/10.1002/smj.2614>
- Stark, H. (2017). As robots rise, how artificial intelligence will impact jobs. *Forbes*. <https://www.forbes.com/sites/haroldstark/2017/04/28/as-robots-rise-how-artificial-intelligence-will-impact-jobs/>
- Tippins, M. J., & Sohi, R. S. (2003). IT competency and firm performance: Is organizational learning a missing link? *Strategic Management Journal*, 24(8), 745–761. <https://doi.org/10.1002/smj.337>
- Tolan, S., Pesole, A., Plumed, F., Macias, E., Hernandez-Orallo, J., & Gómez, E. (2020). Measuring the occupational impact of AI: Tasks, Cognitive Abilities and AI Benchmarks. Working Paper.
- Tomaskovic-Devey, D., Zimmer, C., Stainback, K., Robinson, C., Taylor, T., & McTague, T. (2016). Documenting desegregation: Segregation in American workplaces by race, ethnicity, and sex, 1966–2003. *American Sociological Review*, 71(4), 565–588. <https://doi.org/10.1177/000312240607100403>
- Utterback, J. M. (1987). Innovation and industrial evolution in manufacturing industries. In B. R. Guile & H. Brooks (Eds.), *Technology and global industry: Companies and nations in the world economy*. National Academy Press.
- Vakili, K., & Kaplan, S. (2019). *Organizing for innovation: How team configurations vary with modularity and breadth of application (SSRN Scholarly Paper ID 3509461)*. Social Science Research Network. <https://doi.org/10.2139/ssrn.3509461>
- Van Reenen, J. (2018). Increasing differences between firms: Market power and the macro-economy. In *CEP discussion papers (No. dp1576; CEP Discussion Papers)*. Centre for Economic Performance, LSE. <https://ideas.repec.org/p/cep/cepdps/dp1576.html>
- Wacker, J. (2017). How much will AI decrease the need for human labor? *Forbes*. <https://www.forbes.com/sites/quora/2017/01/18/how-much-will-ai-decrease-the-need-for-human-labor/>
- Webb, M. (2020). The impact of artificial intelligence on the labor market. Working Paper.

Wuebker, R., Saouma, R., & McGahan, A. (2018). Strategy in the age of intelligent machines. *Academy of Management Global Proceedings, Surrey (2018)*, 5. <https://doi.org/10.5465/amgbproc.surrey.2018.0005.abs>

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

How to cite this article: 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*, 1–23. <https://doi.org/10.1002/smj.3286>

