Displacement or Augmentation? The Effects of AI Innovation on Workforce Dynamics and Firm Value

Mark A. Chen[1], Joanna Wang[2]

December 2024

Abstract

This paper studies the effects of Artificial Intelligence (AI) innovation on firm-level employment dynamics and corporate valuation. Applying state-of-the-art large language models (LLMs) and Generative AI to U.S. patent data during 2007-2023, we identify AI-related innovations in seven key functional areas. Using microdata on individual workers' skills and job transitions, we find that AI innovations related to engagement, learning, or creativity augment human labor, but those related to perception displace it. Augmenting AI innovations raise firm-level productivity, while displacing AI innovations lower operating costs. We also find that augmenting (displacing) AI innovations yield more (less) positive valuation effects when the innovating firm has better access to prospective hires (higher costs of terminating employees). Overall, our findings suggest that AI innovations can bring large potential value gains to innovating firms, but how much of those gains are realized depends critically on what frictions are present in the external labor market.

JEL Classification Codes: G30, G32, O32, O33

Keywords: Artificial Intelligence (AI), firm innovation, patents, labor augmentation, labor displacement, generative AI, large language models (LLMs), firm valuation

---

* We are grateful to Tania Babina, Will Cong, Sabrina Howell, Allen Hu, Cen Ling, Xing Liu, Gregor Schubert, and participants at the Columbia & RFS AI in Finance Conference, the DEFT Academy/Xueshuo Summer Institute in Digital Finance, and the 6th PHBS-CUHKSZ Economics and Finance Workshop for valuable feedback. We also thank Jo Suddreth for providing capable research assistance.
[1]Corresponding author. J. Mack Robinson College of Business, Georgia State University, Atlanta, GA 30303, USA. Email: machen@gsu.edu; Tel.: 404-413-7339.
[2] PHBS Business School, Peking University, Shenzhen, 518055, China.

1. Introduction

The rapid development of artificial intelligence (AI) technologies over the past few years has led to growing interest among academics, practitioners, and policymakers in understanding the possible effects of systems that can process vast amounts of information, generate predictions and original content, and engage with human users. Many view AI technologies as a key driver of future economic growth (Aghion et al., 2017; Agrawal et al., 2019a; Furman and Seamans, 2019), with a large part of the gains coming from AI's ability to complement and augment the productivity of knowledge workers.[1] At the same time, questions have arisen about the implications of AI technology for the labor market. As increasingly powerful AI systems begin to perform cognitive tasks that have traditionally been within the domain of skilled human labor, concerns have grown that widespread use of AI may lead to the displacement of skilled knowledge workers (Frey and Osborne, 2017), who have traditionally been considered as the beneficiaries of the spread of older automation technologies.[2]

Does AI technology complement human workers in cognitive tasks, thereby augmenting workforce productivity and creating new jobs and new demand for skilled labor? Or is it the case that AI can perform workplace tasks more inexpensively and more efficiently than humans can, leading to skill obsolescence, labor displacement, and higher unemployment in the future? Despite the obvious economic and public-policy importance of such questions, they have been difficult to address due to a lack of large-scale evidence on AI innovation and workforce dynamics (Seamans and Raj, 2018). In this paper, we use textual data from millions of patents over 2007-2023 and microdata on worker flows to provide some of the first causal evidence on how AI innovation impacts firm-level employment, efficiency, and value.

As a general-purpose technology (GPT), AI includes a remarkably broad range of components, functionalities, and real-world applications (Brynjolfsson and McAfee, 2017; Caliskan and Lum, 2024). Thus, empirically studying the effects of AI is challenging because different types of AI technology may have markedly different effects on firms and labor markets. To address this practical challenge, we propose a new approach to studying AI in terms of its

[1] See e.g. Goldman Sachs (2023), which estimates that advances in AI, such as ChatGPT and other generative AI systems, can complement most knowledge workers' jobs and potentially lead to a 7% rise in global GDP over a 10-year period. Academic studies that highlight AI's ability to complement and improve workers' performance include, for example, Brynjolfsson et al. (2018), Ernst et al. (2019), Webb (2019), and Lane and Saint-Martin (2021).

[2] Reflecting such concerns, the U.S. President's Executive Order issued October 30, 2023 cited disruptions to the workforce caused by job displacement as one of the key potential risks posed by AI (United States, 2023).

functional capabilities. Specifically, we employ textual data from millions of U.S. patents over 2007-2023 to identify AI technologies in areas that correspond to key aspects of human intelligence. We apply Generative AI and Large Language Models (LLMs) to categorize AI innovations into seven overlapping functional areas: language, perception, engagement, inference, decision-making, learning, and creativity.[3]

We document a number of key facts about the growth of AI innovation across functional areas, technological areas, and industrial sectors. For instance, we find that, compared to all patenting in the U.S., AI patenting has grown extremely rapidly over the past decade. Although the largest source of the more than 140,000 AI innovations in our sample is firms in the Manufacturing sector, a substantial amount also occurs across other industries such as Transport, Storage, and Communications. The most common types of AI innovation in our sample are engagement-based AI and creativity-based AI. Creativity-based AI is the single fastest-growing category, likely reflecting the tremendous growth in data analysis and machine-learning algorithms in recent years.

In the main part of our analysis, we examine whether the different categories of AI innovation have different effects on firms' employment with respect to different occupations. For this purpose, we use a new dataset that details the job transitions of individual workers to, from, and within firms in our sample. These data, which are obtained from Revelio Labs, cover employee data at nearly all U.S. publicly-traded firms (including their subsidiaries) from 2008 to the present. In addition to detailing worker movements at a monthly frequency, the dataset also has information on employees' occupations and specific within-firm roles, enabling us to examine how individual occupational segments within firms are affected by different types of AI innovation.

We use data from O*NET on the skill requirements of different occupations to construct measures of the extent to which each occupation is exposed to each type of AI technology. By doing so, we can distinguish between employment effects that are likely the result of direct complementarities or substitutions between AI systems and human workers versus those effects that merely reflect indirect employment effects, between-occupation spillovers, or shifts in the firm's labor demand due to a re-prioritization of various occupations.

In studying the link between AI innovation and firm-level employment, one important issue that arises is potential endogeneity: observed relations between innovation and employment could

[3] Section 2 discusses the rationale for this categorization and provides definitions and examples for each category.

reflect non-causal correlations arising from the influence of other factors. For example, AI patenting activity might be correlated with changes in labor markets, industries, consumer demand, or corporate policies (e.g., overall patenting or R&D spending) that are the true drivers of changes in firms' labor-market strategies. To address these concerns, we follow the approach of prior work that uses the quasi-random assignment of patent applications to patent examiners within different U.S. Patent and Trademark Office (USPTO) art units (Gaule, 2018; Sampat and Williams, 2019; Farre-Mensa, Hegde, and Ljungqvist, 2020). Given that some patent examiners are more lenient than others, the quasi-random assignment of patent applications induces exogenous variation in the probability that a firm successfully obtains, for example, a certain number of AI patent grants within a certain category.[4]

Using employment changes at AI-exposed occupational segments within firms, we document evidence that AI innovations have both augmenting and displacing effects. In particular, engagement-, learning-, and creativity-based AI innovations significantly increase employment at occupations exposed to these types of AI. Perception-based innovations, in contrast, significantly decrease employment at occupations exposed to perception AI. These results are robust to instrumental variables (two-stage least squares) regression analysis as well as to the use of firm fixed effects, fixed effects for year-by-application count totals for different AI types, industry-by-year fixed effects, and a variety of time-varying firm-level controls. Additional results reveal that there are two distinct forms of augmentation: a net increase in external hires who bring new skills to an occupation ("scope augmentation") and a net increase in hires who do not bring new skills ("core augmentation").

Next, to study how AI innovation affects firm-level outcomes, we group the various types of AI patents into four broad categories based on the findings from the occupation-level regressions. We define a patent to be *augmenting* if it belongs to any of the AI categories that are observed to increase overall employment within a firm's occupational segment: creativity, engagement, learning, language, or decision-making. A patent is defined as *scope-augmenting* if it is a type of AI that increases employment of workers who bring new skills to a firm's occupational segment: engagement, learning, or creativity. A patent is *core-augmenting* if it is augmenting but not scope-

[4] Some other papers that use quasi-random assignment of patent applications to capture exogenous variation in patent grants include Melero et al. (2020), and Yang and Yuan (2022).

augmenting. Finally, a patent is *displacing* if it leads to a reduction in occupational employment, i.e., if it is perception-related AI.

We use regression analysis to explore the causal effects of these broad groups of AI innovation on innovating firms' productivity, costs, and value. A priori, there are several reasons to expect that both augmenting and displacing AI patents can yield benefits to innovators. For instance, new AI tools that enhance the productivity of existing workers in a firm (Ersahin, 2020; Yang, 2022) could prompt the firm to scale up production by hiring more workers of the same type. Alternatively, an AI innovation could create value by fostering entirely new jobs (Bessen, 2018; Badet, 2021) or tasks (Acemoglu and Restrepo, 2018) that themselves require additional hiring or reskilling of workers with the firm. AI could also lead to increased firm growth and product innovation (Babina et al., 2024). In the case of labor-displacing AI, a firm could potentially reap large cost savings by laying off human workers and substituting for them with automated AI that is equally capable, but less expensive (Hussey, 2003; Seamans and Raj, 2018).

We find evidence from OLS regressions and instrumental variables regressions that both augmenting and displacing AI patents yield positive benefits to innovating firms, but they do so through different channels. In particular, augmenting innovations increase total factor productivity for the innovating firm. Displacing innovations, however, do not appear to lead to productivity gains for the firm. Instead, displacing AI significantly reduces the firm's operating and wage costs. With regards to corporate valuations, all four broad groupings of AI innovation types cause increases in firms' Tobin's Q in the year following firms' patenting activity.

In the last part of the analysis, we take up a simple empirical question: when are the value gains from labor-augmenting or labor-displacing AI the largest? We conjecture that augmenting patents contribute the most to firm value when the costs of external search, recruiting, and hiring from the external labor market are low. For displacing AI, firms should derive more value when termination costs borne by the innovator are low. To examine these conjectures, we estimate regressions of firm value on augmenting and displacing AI innovation using subsamples based on cross-sectional differences in characteristics such as the enforceability of Non-compete Clauses (NCCs), unemployment insurance coverage, accessibility and supply of potential workers, and the degree of skill specificity to which a firm is exposed through its industry and geography. In general, we find that when firms face greater frictions and impediments to external hiring, the value impact of augmenting AI innovations is less positive or insignificant. For labor-displacing AI, value gains are greater with stronger state-level enforceability of NCCs and lower unemployment insurance coverage. These findings are broadly consistent with the view that labor-market frictions matter in important ways for both innovation-induced external hiring and worker terminations.

Our paper contributes to several strands of literature, including research in finance and economics pertaining to AI innovation, firm-level employment, and the application of text analytics. First, we add to a stream of papers that study the adoption of AI technology and its consequences for certain occupations. Aghion et al. (2017) and Agrawal et al. (2019b) provide a conceptual framework for the effects of AI on job task displacement and augmentation. Some empirical studies examine the effects of AI and big data technologies on financial analysts (Grennan and Michaely, 2020; Cao et al., 2024; Abis and Veldkamp, 2024) and the effects of robo-advisors on financial advisors and asset management (Rossi and Utkus, 2021; Kumar, 2023). By using data on individual workers' job transitions, we can provide a more nuanced and complete picture of how AI affects employment across a broad range of occupations.

Our work is also directly related to recent papers that explore the implications of AI technologies for firm growth and employment. Specifically, Babina et al. (2024) find that AI-investing firms experience higher employment growth and market valuations through increased product innovation. Eisfeldt et al. (2024) find that firms with high exposure to generative AI (i.e., the public release of ChatGPT) experience less frequent hiring and higher stock returns on account of labor-technology substitution.[5] To date, however, researchers have not distinguished among different types of AI innovations and their particular effects on different occupations. By using individual-level employment data, patent text data, and generative AI techniques, we are able to fill this gap with novel empirical evidence. Our functional approach that categorizes AI innovation into seven different types enables us to more precisely examine how the different innovation types impact labor demand and how firm-level productivity and value respond to the arrival of new AI technologies.

Third, our paper contributes to the literature that examines the broad economic implications of automation, skill-biased technological innovation, and creative destruction. Several existing

[5] Alderucci et al. (2020) also document a positive association between AI-related inventions and employment growth. In contrast, Huang et al. (2023) find that freelancers in various occupations experience reductions in both employment and earnings after the mainstream arrival of generative AI. In another recent paper, Acemoglu et al. (2022b) find that employers who adopt AI simultaneously reduce hiring in non-AI positions and alter the skill requirements of remaining postings.

papers focus on the consequences of automation technologies, narrowly defined. For instance, researchers have used the stock of industrial robots in different industries and countries (Acemoglu and Restrepo, 2020; Bates et al., 2021; Qiu et al., 2021) as a proxy for the extent of physical automation. Some recent studies emphasize that technological innovation may create new jobs and reinforce labor growth (Acemoglu and Restrepo, 2019; Autor and Salomons, 2018). Other papers discuss the skilled-biased feature of technological innovation (Autor et al., 2003; Acemoglu and Autor, 2011) which affects high and low-skilled labor in different ways. Griliches (1969) and Krusell et al. (2000) argue that equipment and information technology (IT) capital are more complementary to skilled labor than to unskilled labor. To study the effects of AI, Kogan et al. (2023) develop measures of labor-saving and labor-augmenting technology. Motivated by such work, our paper focuses on AI, widely considered to be one of the broadest and most impactful general-purpose technologies to have emerged in recent years. However, rather than pre-defining labor-augmenting or displacing technologies, we take an agnostic approach and study which specific types of AI innovations are truly augmenting or displacing with respect to different occupations. Thus, we can obtain a much more complete picture of the varied scope of AI technologies and their consequences for firm-level outcomes.

Finally, our work offers a new methodological approach to the study of large-sample textual data that can potentially be applied to many other questions in finance and economics. A large and growing literature in financial economics uses Natural Language Processing (NLP) techniques to study corporate filings, disclosures, patent texts, and other textual data. Earlier studies in this area use text-filtering or word-embedding algorithms that are easier to interpret, but less accurate (Loughran and McDonald, 2011; Henry and Leone, 2016; Routledge, Sacchetto, and Smith, 2013; Webb, 2019). More recent studies exploit techniques related to machine-learning and generative AI (Chen et al., 2019; Giczy et al., 2022; Li et al., 2024, Jha et al., 2024) to analyze textual data. We build on this line of research by applying state-of-the-art Generative AI techniques to classify AI innovations at scale with less ambiguity than would be possible with traditional text filtering or earlier machine-learning methods.

The remainder of the paper is organized as follows. In Section 2, we discuss our identification of seven AI categories. Section 3 describes our data sources and provides descriptive statistics. Section 4 describes our methodology on AI technology-skill mapping, occupational AI exposure, and the instrumental variables approach to capturing exogenous variation in AI innovation. In Section 5, we report our results. Section 6 concludes.

2. A Functional Categorization of AI

While there is no single, universally accepted definition of AI, a view commonly held by researchers and practitioners is that AI consists of computing systems that can accomplish tasks that have traditionally required human intelligence. The idea that artificial intelligence can or should do what human intelligence is capable of has featured prominently in the work of computer scientists since at least the mid-20th century. For example, in 1950 Alan Turing proposed a test based on an "imitation game" whereby a computer can be said to possess artificial intelligence if it can adequately mimic human responses under certain conditions. (Turing, 1950). In 1956, John McCarthy, often described as the "Father of AI," organized a summer conference that is widely regarded as the birth of AI as a scientific field. The conference proposed to study machines based on the conjecture that "every aspect of learning or any other feature of [human] intelligence can in principle be so precisely described that a machine can be made to simulate it."

In the decades since the founding of the AI field, major technical advances such as artificial neural networks, deep learning, and generative AI have enabled artificial systems to equal or exceed human performance in specific, narrowly-defined cognitive tasks. Many computer scientists and engineers believe that progress will eventually lead to Artificial General Intelligence (AGI), which is a form of machine intelligence that, even without extensive pre-training, would be able to think, reason, create, and generalize as humans do. While AGI currently remains only a hypothetical possibility, the development of intelligent systems that can fully match humans across the broad range of their cognitive tasks persists as a central goal of research in the field.

Motivated by the oft-discussed parallels between AI and human intelligence, we propose to categorize AI innovations according to their key cognitive capabilities. Our functional typology builds upon two well-known organizing frameworks that have been widely used in both theory and practice to study human cognition. First, we draw upon the unified theory of cognition due to Allen Newell, a prominent 20th-century researcher in computer science and cognitive psychology. According to Newell's theory, all cognitive behavior is based on a single set of mechanisms (Newell, 1990). These mechanisms cover six major areas: (1) problem-solving, decision-making, routine action; (2) memory, learning, skill; (3) perception, motor behavior; (4) language; (5) motivation, emotion; and (6) imagining, dreaming, daydreaming.

The second framework that we rely on is the set of key neurocognitive domains defined in the 5th edition of the *Diagnostic and Statistical Manual of Mental Disorders* (DSM-5). Published in 2013 by the American Psychiatric Association, DSM-5 is a standard reference guide used around the world by clinicians, researchers, hospitals, pharmaceutical companies, and policymakers. DSM-5 specifies six principal neurocognitive domains: (1) complex attention; (2) executive function; (3) learning and memory; (4) language; (5) perceptual-motor function; and (6) social cognition. Each of the principal domains in DSM-5 also contains subdomains that help to characterize specific types of neurocognitive functioning within the broader domain.

Table 1 gives a side-by-side summary of the key cognitive domains specified by Newell (1990) and those specified by DSM-5. As seen in the table, some of the domains correspond well between the two frameworks. However, the "complex attention" domain in DSM-5 has no obvious mapping to a single area in Newell's theory, and Newell's "imagining, dreaming, and daydreaming" has no single counterpart in DSM-5. We therefore harmonize the two frameworks as follows. First, we associate both "complex attention" and "executive function" with Newell's "problem-solving, decision-making, routine action" category. This view is consistent with an extensive body of research in psychology and cognitive sciences showing that complex attention plays an indispensable role in problem-solving and higher mental processes. (e.g., Smith and Ratcliff, 2009; Knippenberg et al., 2015). Second, in line with the argument by many cognitive scientists, psychologists and philosophers that imagination and memory are closely related,[6] we map Newell's areas of "memory, learning, skill" and "imagining, dreaming, and daydreaming" to the single DSM-5 domain of "learning and memory."

Based on our harmonization of the two organizing frameworks, we propose a novel typology for studying AI. This typology categorizes AI systems according to seven different cognitive capabilities, which we list in the third column of Table 1. These capabilities include *inference*, *decision-making*, *learning*, *creativity*, *language*, *perceptual-motor*, and *engagement*. Each capability stems from one of the principal cognitive domains and represents a leading

[6] According to philosopher Thomas Hobbes, "Imagination and Memory are but one thing…." (Hobbes, 1968). Neuroimaging evidence also indicates that memory and imagination are cognitive functions that are very closely related (see, e.g., Mullally and Maguire, 2014).

example of how an AI system might function within that domain. We describe the seven AI capabilities as follows:

- ***Language***: Understanding and using language through speaking, listening, reading, or writing
- ***Perception/Motor***: Gathering, organizing, and identifying sensory information from external stimuli to gain an awareness of the environment
- ***Engagement****:* Understanding, interacting with, and responding to the behavior of human users
- ***Inference****:* Drawing conclusions based on evidence, logical reasoning, or heuristic reasoning
- ***Decision-Making****:* Evaluating options based on relevant information and subsequently choosing an alternative or a course of action
- ***Learning****:* Acquiring new knowledge, skills, values, preferences, or behaviors
- ***Creativity****:* Transcending traditional ideas, rules, patterns, or relationships to generate new ideas, forms, methods, interpretations, or possibilities

Table 2 lists some specific tasks/processes and sample use cases associated with each AI capability. It should be noted that, in our typology, an AI system need not be associated with just one capability. A conversational AI agent, for example, might demonstrate multiple functions, including an ability to engage with users, inference ability, perceptual ability, and creativity. We also note that the categories in our typology are intended to capture the degree to which technologies function as part of the intended AI system. Thus, for instance, an improved design for a generic microprocessor would not be considered an AI innovation or technology even though one of the possible uses of the microprocessor is as a hardware component of an AI system.

3. Data

3.1 Sample

Our source of data on firms' innovation is the full-text database of patent applications and grants at the USPTO Bulk Data Storage System (BDSS)[7]. From the PatentsView website[8], we obtain information on individual patents' citations and their International Patent Classification

[7] https://bulkdata.uspto.gov/
[8] www.patentsview.org

(IPC) codes. To identify the patenting activities of U.S. public firms, we proceed to merge together the USPTO data with CRSP/Compustat. We start by building links within the USPTO data between patent applications and patent grants. From Jan 2007 to December 2023, there are 5,913,659 utility patent applications, and 3,765,176 of them are observed to have been granted. Next, from the data sample[9] of Kogan et al. (2017), henceforth, KPSS, we obtain the linking data between USPTO patent grant numbers and CRSP PERMNO (PERMCO) identifiers. (This sample covers 3,160,453 patents from 1926 to 2022). To identify the patent application information up to 2022, we retrieve the internal link between grants and applications from USPTO and match it with the KPSS data. Among patents for which both application and grant information are available, we have 1,014,772 patents filed by CRSP U.S. firms from the KPSS data.

The previous step does not include (1) patent applications that were not yet (never) granted from 2000 to 2022; (2) patent applications that were not yet (never) granted in 2023; and (3) patent grants in 2023. Thus, we use a basic name-matching algorithm to supplement the linking between CRSP firms and USPTO patents. From the KPSS data, we use PERMNO (PERMCO) identifiers to retrieve all relevant firms' historical names from the CRSP database. For patent grants, we directly downloaded the assignee data from the Patents View database. [10] We then extract patent assignee (applicant) information for all applications from the USPTO bulk database.[11]Using the name mapping between granted patent assignee and CRSP firm historic names, we further identify 961,655 patent applications before 2022 as well as 45,123 patent applications and 21,132 grants in 2023.[12] Finally, we obtain other firm-level data from the CRSP/Compustat merged database. The resulting sample consists of 2,307 unique U.S. firms that file 1,562,529 patent applications from 2007-2023, of which 1,253,239 are granted by the end of 2023.

3.2 Using Generative AI to identify and categorize AI patents

Until recently, the standard approach in finance and economics for analyzing textual data has been to use text filtering or word embedding techniques to classify documents according to

---

[9] This dataset provides an updated data series for KPSS (Kogan, et al., 2017) values and forward citations, a patent number to the CRSP "PERMNO" match, and a patent number-to-CPC class match following the paper.

[10] The Patents View database (https://patentsview.org/download/data-download-tables) only provides assignee information for granted patents.

[11] For patent applications, we cleaned and parsed the USPTO patent application XML files and extracted patent assignee names if they were recognized as organizations.

[12] Specifically, these additional patents are identified when their assignee names can be found in the KPSS matched pairs between CRSP firm names and USPTO assignee names.

single-word frequencies or the co-occurrence of certain pairs of words.[13] While such methods are straightforward and easy to interpret, they do not make full use of the semantic connections between words, sentences, and documents that can dramatically improve textual analysis. The recent development of generative AI tools such as ChatGPT and Google Gemini has made it possible to retrieve, categorize, and analyze documents using simple queries that refer to general topics or concepts. These instances of Large Language Models (LLMs) are trained on trillions of tokens and can respond to queries using a large set of commonsense knowledge.[14]

In view of the difficulties with using traditional NLP techniques to capture the rich semantic content in patent filings, we turn to generative AI to help identify and classify different types of AI innovations. Specifically, we use a near state-of-the-art LLM called *Qwen2.5-14b-Instruct*. This generative AI model was developed by Alibaba Cloud and released to the public in September 2024 under an open-source license for broad commercial and research use. The model features 14.8 billion parameters and is an instruction-tuned version of a base model that was trained on data consisting of approximately 18 trillion tokens (https://github.com/QwenLM/Qwen2.5). According to several performance benchmarks, as of October 2024, *Qwen2.5-14b-Instruct* is the top-performing LLM in its size class and outperforms many larger, proprietary models. For instance, within the topical area of Engineering, *Qwen2.5-14b-Instruct* is seen to outperform well-known models such as Gemini-1.5-Pro, Claude 3-Opus, and GPT-4 Turbo.[15]

Our overall approach to identifying and categorizing AI innovations from among the millions of U.S. patent applications filed during 2007-2023 consists of two main steps. First, for each patent application, we formulate a query to elicit from the LLM an inference about whether the invention is AI-related. To reduce ambiguity and to obtain from the LLM as clear a response as possible, we use prompts that feature a specific yes/no question combined with the text of a patent abstract. For example, to check whether or not U.S. Patent Application #20210209449 is related to artificial intelligence, we prompt the LLM with the following question/instruction, along with the patent

[13] Some notable examples of this approach to computational linguistics include Tetlock et al. (2008), Hanley and Hoberg (2010), Loughran and McDonald (2011), Jegadeesh and Wu (2013), Hoberg and Phillips (2016), and Li et al. (2021).
[14] An emerging literature in finance and economics applies LLM and generative AI techniques for text classification and content analysis. Papers in this vein include Bartik et al. (2023), Chang et al. (2024), and Li et al. (2024), among others. Eisfeldt and Schubert (2024) provide a review of recent finance research that uses generative AI techniques to analyze and classify textual data.
[15] See, e.g., the LLM performance benchmarking leaderboard at https://huggingface.co/spaces/TIGER-Lab/MMLU-Pro (accessed October 18, 2024).

abstract as context: "*Is the following invention directly related to artificial intelligence? Respond with just YES or NO. Abstract: Graphical elements in a user interface (UI) may be detected in robotic process automation (RPA) using convolutional neural networks (CNNs). Such processes may be particularly well-suited for detecting ... [remaining abstract text follows]".* Note that we do not require the LLM to provide an explanation for its yes/no inference; doing so would greatly extend the time required for each inference.[16] Applying this querying procedure to all 5, 323, 321 patent applications in the sample yields a classification of each patent application as being either AI-related or not.

Next, for each patent application that is AI-related, we submit seven queries to the LLM to ascertain whether or not the application belongs to each of the seven AI categories. We again use simple yes/no prompts to ensure speed of inference. For instance, to check whether a patent application is related to perception-based AI, we query the LLM with the following prompt: "*Is the following invention directly related to artificial intelligence that has perceptual ability? Respond with just YES or NO. [patent abstract text follows]*". The overall collection of binary yes/no LLM responses constitutes a full categorization of AI patent applications in our sample.

To check whether our LLM-based approach is able to identify and distinguish between different types of AI innovations with a high degree of reliability, we further conduct in-depth queries for each patent application in a random subset of the AI-related applications. The queries we use for this purpose are similar to the queries used for the entire sample, except that they additionally request the LLM to provide an explanation for its yes/no inference. For instance, to understand whether and why the LLM considers a given AI patent application to be perception-based, we use the prompt "*Is the following invention directly related to artificial intelligence that has perceptual ability? Respond with YES or NO and give an explanation. [followed by the patent abstract text]".*

Appendix A provides examples of yes/no classifications for each AI category along with accompanying explanations from the LLM. Based on a review of these examples and numerous others, we conclude that the LLM is capable of differentiating between fine shades of meaning in patent abstracts. For instance, consider the case of Patent Application # 20220058743, "User

[16] It is well-known that token generation accounts for a large majority of the processing time required for an LLM to respond to a query. Our approach of eliciting simple yes/no responses from an LLM without requiring explanations has the benefit of dramatically shortening the overall time required for inference. Moreover, as discussed below, directing the LLM to forego providing explanations does not seem to compromise the quality of the LLM's inferences.

Interactions in Mobile Damage Assessment and Claims Processing," (Allstate Insurance Company), listed in Appendix A. Despite the fact that the abstract for this patent filing does not contain the term "perception" or any other synonymous terms, the LLM correctly recognizes that the invention demonstrates AI and visual perception ability because the disclosed system can automatically collect and analyze information from images, photos, and video. Also, we note that the LLM is able to reliably determine when an AI invention does not relate to a particular category. For example, the LLM correctly infers that Patent Application #20230162412, filed by Siemens Medical Solutions USA, Inc., does not directly pertain to decision-making (see Appendix A). As explained by the LLM, although this invention uses an artificial neural network to reconstruct 3D images from 2D projections, its capabilities relate to pattern recognition and data transformation rather than making actual decisions based on the patterns and data.

3.3 Data on firm-level, by-occupation employment

We obtain data on occupation-specific employment at U.S. firms from a novel database compiled by Revelio Labs. This database aggregates global workforce dynamics of public employer and employee records from more than 4.5 million companies and 1.1 billion resume profiles, recording individual-level data on demographics (e.g., name, gender, race, origin), job position (affiliated firm identify, title, job role[17], salary, tenure), skill taxonomy, education, and job transition.[18] In addition, the 8-digit SOC occupation code (from O*NET database) for each individual is also provided in the data. Our analysis of firms' workforce dynamics relies on three key types of employee-level data. First, data on each employee's occupation and employment record enable us to construct measures of occupational employment changes. For instance, using the time-varying distribution of individual workers in each occupation-firm-year, we can calculate the firm-occupation-level yearly labor growth rate for workers in different occupations. Second, by linking the required skills of each occupation from the O*NET database with seven categories

[17] The job role taxonomy is clustered by mathematical representations of each job using the title, the text description of the position (from either individual describing their own experiences or employers on a job posting), and individuals' skills, associates, and previous experience. The role taxonomy is adjusted periodically to adapt to the changing occupational landscape.

[18] Specifically, the job transitions data give us a full picture of an employee's work history, current status, and job transfers within or across firms over time. The transition data also allow us to track detailed information on characteristics of new (previous) firms and new (previous) jobs, including geographic location, firm identity, role name, salary, and job start (end) date. The job role taxonomy is clustered at different levels of granularity. In this paper, we use job roles with 1000 clustering.

of AI technologies based on their functional capabilities, we can identify occupations that are exposed to AI. This enables us to detect the direct substitutive and complementary effects of AI on human workers who perform tasks in occupations exposed to AI. Third, we obtain individual-level skills data (with a full set of 3,000 different skills) from online profiles recorded by Revelio Labs. We combine the individual user skill data and transition data to identify the skill novelty of newly hired and separated workers. This helps us characterize firms' skill accumulation process in response to AI innovation. Specifically, we are able to distinguish between scope augmentation, whereby AI complements human workers with new skills, and core augmentation, whereby AI complements human workers with existing skill sets.

3.4 Variable construction

3.4.1 Occupational employment variables

To examine how AI innovation affects different types of human workers, we construct two outcome variables using the Revelio Labs individual position file. First, for each 3-digit SOC O*NET occupation and each firm-year, we measure the occupational-level *Employment growth (in headcount)* as the yearly change in headcount scaled by the firm's total employees. We then merge the individual transition data with the individual skill data to identify newly hired (fired) workers with unique skills compared to the median of their occupations. Specifically, for each individual, we construct his/her skill set vector based on a total of 3,000 possible skills in the data. We then perform an element-wise comparison between each individual's skill vector and the median skill vector (computed across workers in the occupation). We define workers with new skills to be those with an individual skill vector that contains any skill elements missing from the occupational median worker's skill set. Based on this, we measure another occupational-level outcome, *Employment growth (with new skills),* as the yearly changes in number of employees with new skills for each occupation, scaled by the firm's total employees. Table B2 in Appendix B provides examples of individual skill sets for both hired and separated employees. Our later analysis makes use of data on the skills of transitioning workers, which allows us to examine a form of AI labor-augmentation that depends on adding new skills to an existing.

3.4.2 Firm-level AI innovation variables

To study the effects of AI innovation in each category on occupational employment growth, we count the number of patent grants in seven AI categories defined in Section 2, *Log # Type AI patents*, and all other patent grants, *Log # Non-type AI patents*, as explanatory variables. These patent count variables are instrumented by the total examiner leniency as described in Section 4.2. Later, we focus on the overall augmenting and displacing effects of AI based on their predicted employment outcomes. We re-group the patents and construct the *Log # augmenting (displacing) AI patents* as the log-transformed total number of patent grants in AI categories that are estimated to have positive (negative) relation with the employment growth in occupations that exposed to AI (Table 5 Panel A). Then, we construct the *Log # Scope (Core)-augmenting AI patents* as the log-transformed total number of patent grants in AI categories that are estimated to have positive (insignificant) relation with the employment growth with new skills in exposed occupations (Table 5 Panel B). Additional details for patent re-grouping are reported in Section 5.2.1.

3.4.3 Firm-level outcome variables and control variables

We estimate firm-level total factor productivity as a measure of firm efficiency. As discussed in the literature (Nadiri, 1970; Craig and Harris, 1973; Bennett, Stulz, and Wang, 2020), TFP increases as a firm uses its inputs (i.e., capital and labor) effectively. We obtain the firm-level TFP from the residual estimates in the firm-year level linear production function model with capital and labor inputs, controlling for firm fixed effects. Firm-level production inputs from capital and labor are separately proxied by cost of cost good (from Compustat) and firm total employment (from Revelio Labs). The measure of the firm production output is proxied by net sales (from Compustat). Then, the dependent variable in the firm-level test, $\textit{Total Factor Productivity}_{t,\ t+1}$, is the log-transformed changes in the yearly TFP. The residual estimates of TFP are scaled in thousands before the log transformation. We use two measures of costs related to firm production and operation. The dependent variables are the growth rate in percentage for total operating expenses and SG&A expenditures from t to t+1. Both total operating expenses and SG&A expenditures are from Compustat.

Last, we construct a measure of firm value using the average Tobin's Q value for the two years after the year of patent filings. Following previous literature (Fazzari et al., 1988; Erickson and Whited, 2012), we calculate firm-year level Tobin`s Q as the ratio of the book value of debt (Compustat items *DLTT* + *DLC*) plus the market value of equity (Compustat items *PRCC_F* ×

CSHO) minus the firm's current assets (Compustat item *ACT*) to the book value of property, plant, and equipment (Compustat item *PPEGT*). For all regressions, we include a set of firm-level control variables constructed using the data from Compustat. *Size* is the natural log of total assets in the prior year from Compustat. *MTB* is the market-to-book ratio in the prior year. *ROA* is income before extraordinary items divided by total assets in the prior year. *R&D* is the natural log of R&D expenditures in the prior year. Missing values of R&D are imputed as zero and indicated by *R&D Missing*.

3.5 Summary statistics

We first report the growth of U.S. AI innovation through patenting activities over time. Figure 1 illustrates the time-series occurrence of AI patenting by U.S. public firms from June 2007 to December 2023. Patent grants are counted as of the publication date of grant. Panel A shows the time series for AI patents and for all patents, while Panel B shows the time series for each AI patent category. Patent frequencies are cumulative and shown on a monthly basis, scaled by the patent count in June 2007 (as the base group). In Panel A, we find a sharper increasing trend of overall AI patenting activities compared to all patenting activities. The growth rate of AI innovation starts to diminish in the year 2021, possibly due to firms' financial stress during the Covid-19 pandemic (Ellul et al., 2020). Panel B shows that AI innovation related to engagement and creativity account for larger fractions of AI innovation compared to other categories. Engagement AI and creativity AI also exhibit higher growth rates in later years during the sample period, likely reflecting the recent mainstream adoption of generative AI tools.

In Figure 2, we report the distribution of AI innovation activity by firm size, labor intensity, and R&D intensity. Panel A reports, by firm size, the firm-level AI patents in each category as a fraction of all of a firm's patents. It implies that larger firms generate more innovation output (2% for all AI patents) about AI technologies than small firms do. Panel B summarizes AI patent shares in labor-intensive and capital-intensive (non-labor-intensive) firms. We find that capital-intensive firms file relatively more AI patents than labor-intensive firms. The difference is mainly driven by decision-making AI and inference AI. The fact that there is less AI innovation in labor-intensive firms suggests that labor displacement is unlikely to be the only driver of firms' AI adoption. Panel C compares AI innovation activity in R&D-intensive and non-R&D-intensive firms. The figure shows that the fraction of AI patents in the two groups exhibit relatively small differences, although there is a larger fraction of all patents in R&D-intensive firms. The evidence suggests that AI innovation does not necessarily rely heavily on R&D as an input into the innovation process.

To gain a better understanding of how AI affects different jobs, we construct heatmaps for the estimates of AI effects on employment growth at the (2-digit SOC) occupation level in Figure 3. For each AI category, we estimate the AI effects by running 85 regressions for each 3-digit SOC occupation on its employment growth in headcounts in a given firm-year. We assign each occupation a score based on the sign of the estimates (significant positive estimates as 1, negative estimates as -1, and insignificant estimates as 0). Then we aggregate the assigned score from 3-digit SOC to 2-digit SOC occupations, weighted by the employment share of each 3-digit SOC occupation. Panel A reports the results for all occupations, and Panel B reports the results for AI exposed occupations as defined in Section 4.1. By unpacking the effects of AI at the occupation level, we observe considerable heterogeneity across occupations and technologies. Specifically, in Panel A, perception AI exhibits large-scale displacement in many occupations, whereas creativity and learning AI augment labor in many occupations. In Panel B, we find a different pattern of AI augmentation and displacement when regressions are restricted to occupations with AI exposure. This suggest that AI innovation may indirectly affect some jobs, likely due to firm-level optimization on labor adjustment.

Next, we report in Table 3 the distribution of AI patent grants across industrial sectors and technology classes based on patent IPC classifications. Industrial sectors that experience the most frequent AI innovation include manufacturing (67,285 AI patent applications), Transport; storage and communication (48,008 AI patent applications), and Electricity; gas and water supply (9,706 AI patent applications). Although innovators in the communication and electricity sectors consist of a large number of high-tech firms, the evidence suggest that AI technologies is still widely adopted beyond high-tech sector. Panel B reports the distribution of AI patent applications and grants in the ten most frequent technology classes based on patent IPC classification. Most categories of patents heavily build on "G&H" technologies, including Computing, Calculating, Counting (G06); Electric Communication Technique (H04) and Musical Instruments and Acoustics (G10)[19].

[19] The Musical Instrument and Acoustics (G10) technology class is not only restricted to musical instruments. It includes technologies applied on sound emitting devices. For example, some artificial intelligence patents on methods or devices for transmitting, conducting, or directing audios can be classified under this category. See more details in https://www.wipo.int/ipc/itos4ipc/ITSupport_and_download_area/20240101/pdf/scheme/full_ipc/en/g10.pdf

Next, we report occupation-level and firm-level characteristics in Table 4. Panel A shows characteristics for occupational employment growth and employee transitions with and without new skills. On average, the occupational employment growth rate on headcount is decreasing, suggesting a shrinking of occupation size over time. However, the growth rate of employees with new skills is positive, whereas the growth rate of employees without new skills is negative. For the decomposition of employment growth into worker inflows and outflows, we find the hiring of new skill workers dominates the hiring of other workers. On the other hand, job separations among employees with new skills are much less frequent than employees without new skills. Overall, the evidence reveals an overall skill bias in the job market. Panel B reports firms' AI innovation and financial characteristics. The firm-level yearly average number of AI patents is 7.07, which is mainly explained by inference AI (4.65) and decision-making AI (4.02). In later sections, we will aggregate across AI technology types and define AI patents as "augmenting" or "displacing" based on estimated firm employment effects (e.g., higher skilled-labor growth) in the year after the patent filings.[20] On average, the frequencies of firms' labor augmenting and displacing AI patents are 4.84 and 2.21, respectively. The average number of AI patent grants that lead to scope (core) augmentation is 3.32 (4.24) for a firm in a given year.

4. Empirical Methodology

Our analysis of whether AI innovation augments and/or displaces human workers requires being able to identify the different impacts that AI innovations tend to have on different occupations. In this section, we describe the key elements behind our empirical approach, including (1) the use of O*NET data on occupational skill requirements to construct occupation-level measures of worker exposure to AI types; and (2) how we use data on USPTO patent examiners to address endogeneity and to study the causal effects of AI innovation on employment changes.

4.1 Measuring exposure of occupations to different types of AI innovation

To construct measures of how related (exposed) a given occupation is to different types of AI innovation, we proceed in several steps. First, we make use of data from the O*NET database (https://www.onetonline.org) on skill requirements underlying each occupation. These skill

[20] A detailed description of how we reclassify patents is given in Section 5.2.1.

requirements are based on a set of 35 skills, including elements such as “Critical Thinking,” “Social Perceptiveness,” “Reading Comprehension,” “Judgment and Decision-Making”, and “Programming.” In the O*NET database, each occupation at the most disaggregated (8-digit) SOC level is associated with numerical importance weightings that indicate how essential each of the 35 required skills is to the occupation. These scores range from 1 to 5, with higher scores indicating greater importance. To simplify the aggregation of skill requirements at the occupation level, we code any skill having a score of 4.0 or above as “important” to the occupation, while any skill with importance below 4.0 is deemed “not important.”

Next, we manually review each of the 35 skills to determine if the skill is fundamentally related to one or more of the seven cognitive areas in our AI typology. Some of the skills link in an obvious way with one or more of the cognitive areas. For instance, the cognitive capability of *language* is clearly related to the O*NET skill of “writing”; while the cognitive capability of *inference* is clearly related to the O*NET skill of “mathematics.” However, in some cases it is not obvious whether a skill links to a cognitive capability. For these cases, we deem a relation to be present if the cognitive capability appears to be a necessary ingredient into actually applying the skill in question.

We then combine the manual mapping between skills and AI types with the O*NET skill importance measures to derive an overall picture for how a given 8-digit occupation relates (via its skill requirements) to the various AI innovations. Thus, suppose for instance that an 8-digit occupation has two skills related to perceptual capability. By construction, each of the two related skills must be not only an important occupational requirement as per O*NET data, but also related to perception-based cognition according to our manual mapping.

In the final step of constructing our exposure measure, we aggregate 8-digit SOC occupation relatedness to the level of 3-digit SOC occupations. Specifically, for a given AI type and a given 3-digit SOC occupation, we check whether at least one 8-digit SOC occupation contained therein is related (via its skill requirements) to the AI type. If so, then we consider the 3-digit occupation to have exposure to the AI innovation type. However, if no sub-occupation is related to an AI type, then the 3-digit occupation is deemed to be unexposed.

Table B1 in Appendix B shows, for each AI type, the top three most-related 3-digit SOC occupations. (Note that, for this ranking, we count the number of related skills for an occupation.) The table shows considerable heterogeneity in the types of occupations that are most related to a given AI type. Nonetheless, some systematic patterns are also evident. For example, Top Executives (SOC Code 111) is the occupation that is most related to Creativity AI, Engagement AI, and Learning AI. Occupations involving supervision are most important for Decision-Making AI, but supervisory occupations are also present among the top 3 for several other AI technological categories.

4.2 Identification

To identify the causal effects of AI patent grants on occupation-level and firm-level outcomes, we exploit exogenous variation in patent approval rates that arises from the randomness in how patent applications are allocated to examiners at the U.S. Patent and Trademark Office (USPTO). This approach to identification, introduced by Gaule (2018), Sampat and Williams (2019), and Farre-Mensa et al. (2020),[21] exploits two key institutional features of the USPTO patent examination process. First, individual patent examiners are observed to differ systematically in their leniency, i.e., in how likely they are to allow a patent application into granted status. Second, although patent applications are sent to different art units (groups of examiners) based on technological attributes, the assignment of patent applications to individual examiners within an art group is largely random (Lemley and Sampat, 2012; Sampat and Williams, 2019).[22]

We build on the approach of Farre-Mensa et al. (2020) to capture exogenous variation in two or more different types of patenting by a firm in a given year.[23] First, for each patent examiner k in each year t, we calculate a time-specific leniency measure, E_{kt}, as the fraction of applications the examiner reviewed during the year that were granted. Because an examiner's art unit could induce a systematic component to his or her leniency, we estimate a simple regression to account for year-specific art unit effects:

$$E_{kt} = \Gamma' \mu_{lt} + \epsilon_{kt}, \tag{1}$$

[21] Gaule (2018) uses variation in USPTO examiner leniency to study the causal effects of patent protection on the success of startup firms. Sampat and Williams (2019) employ a similar approach to explore whether genome patents affect scientific investment and follow-on innovation. In Farre-Mensa et al. (2020), variation in patent examiner leniency is used to study how winning a patent grant causally affects a startup's employment and sales growth.

[22] Lemley and Sampat (2012) conduct written interviews with USPTO examiners about the patent assignment process. They find no evidence of selection based on characteristics of applications other than observed conditions in standard USPTO datasets (i.e., technology type represented by USPTO art unit).

[23] The analysis of Farre-Mensa et al. (2020) only considers a firm's first patent grant and does not distinguish among patent types. By partitioning a firm's patent applications within a year into different groups, we are able to study how different types of patents concurrently influence firm-level outcomes in different ways.

where μ_{lt} denotes art unit-by-year fixed effects. The residual from this regression, denoted by E^*_{kt}, constitutes an adjusted measure of examiner k's leniency in year t that accounts for the overall leniency of the examiner's art unit in the same year. Then, for any given patent category of interest, we can define aggregate leniency measures for patenting activity within and outside of the category:

$$Z_{it,Type} = \sum_{j \in D_{it,Type}} E^*_{k(j,t),t} \quad (2)$$

$$Z_{it,NonType} = \sum_{j \in D_{it,NonType}} E^*_{k(j,t),t} \quad (3)$$

In these definitions, $D_{it,Type}$ ($D_{it,NonType}$) is the set of patent applications filed by firm i in year t within (outside of) the category of interest, and $k(j,t)$ is the random examiner to which patent application j is assigned in year t. In essence, these two equations partition a firm's total innovation activity in a year into two groups (within-category applications and outside-of-category applications) and separately capture how "lucky" the firm is in randomly drawing lenient examiners for the two groups of applications.

With aggregate leniency measures $Z_{it,Type}$ and $Z_{it,NonType}$ in hand, we can then construct instrumental variables that capture exogenous variation in different categories of patenting. For example, to explore the occupation-specific employment effects of patenting within and outside of a given AI category, we first form instruments by interacting aggregate leniency measures with indicators (constructed above) showing whether an occupation is "Exposed" or "Non-Exposed" to the relevant category of AI. Then, pooling together observations at the occupation-firm-year level, we can estimate two-stage least-squares (2SLS) models such as the following:

<u>First stage:</u>

$$\begin{aligned} TypePatents_{it} \times ExposedOcc_{ij} = \ & \alpha_0 + \alpha_1 Z_{it,Type} \times ExposedOcc_{ij} \\ & + \alpha_2 Z_{it,NonTy} \quad \times ExposedOcc_{ij} \\ & + \alpha_3 Z_{it,All} \times NonExposedOcc_{ij} \end{aligned}$$

$$+\alpha_4 X_{it} + \lambda_i + \rho_{kt} + \mu_{kt} + \theta_{it} + \varepsilon_{it} \quad (4)$$

$$\begin{aligned} NonTypePatents_{it} \times ExposedOcc_{ij} = \; & \beta_0 + \beta_1 Z_{it,Type} \times ExposedOcc_{ij} \\ & + \beta_2 Z_{it,No} \qquad \times ExposedOcc_{ij} \\ & + \beta_3 Z_{it,All} \times NonExposedOcc_{ij} \\ & + \beta_4 X_{it} + \lambda_i + \rho_{kt} + \mu_{kt} + \theta_{it} + \nu_{it} \quad (5) \end{aligned}$$

$$\begin{aligned} AllPatents_{it} \times NonExposedOcc_{ij} = \; & \gamma_0 + \gamma_1 Z_{it,Type} \times ExposedOcc_{ij} \\ & + \gamma_2 Z_{it,NonType} \times ExposedOcc_{ij} \\ & + \gamma_3 Z_{it,All} \times NonExposedOcc_{ij} \\ & + \gamma_4 X_{it} + \lambda_i + \rho_{kt} + \mu_{kt} + \theta_{it} + \omega_{it} \quad (6) \end{aligned}$$

Second stage:

$$\begin{aligned} Y_{ij,t+n} = \; & \delta_0 + \delta_1 \widehat{TypePatents_{it} \times ExposedOcc_{ij}} \\ & + \delta_2 \widehat{NonTypePatents_{it} \times ExposedOcc_{ij}} \\ & + \delta_3 \widehat{AllPatents_{it} \times NonExposedOcc_{ij}} \\ & + \delta_4 X_{it} + \lambda_i + \rho_{kt} + \mu_{kt} + \theta_{it} + u_{it} \quad (7) \end{aligned}$$

In the above two-stage model, $Y_{ij,\,t+n}$ is a measure of employment change for occupation j at firm i. $ExposedOcc_{ij}$ ($NonExposedOcc_{ij}$) is an indicator equal to one if occupation j has (does not have) exposure to the AI technology in question. $TypePatents_{it}$ is the log of one plus the number of granted within-category AI patents filed by firm i in year t, and $NonTypePatents_{it}$ is the log of one plus the number of all other granted patents filed by firm i in year t. Following Farre-Mensa et al. (2020), we count a patent application as occurring in year t if the "First Office Action" date from the USPTO is during the year.[24]

[24] The First Office Action date for a patent application is the first date on which the USPTO contacts the applicant by mail about the examiner's initial examination and about material deficiencies, if any, that must be remedied for allowance to a grant to occur. This date corresponds to the earliest date on which the applicant learns of the identity of the examiner (and hence the likelihood that the patent will be granted). For our analyses relating to firm valuation (Tobin's Q), we instead count patents based on the later of the First Office Action date and the application publication date (i.e., the first time at which the public and market participants at large learn about the existence and status of a patent filing).

The vector X_{it} denotes a set of control variables including firm size, firm performance, and R&D expenditures. The term λ_i represents firm fixed effects. The terms ρ_{kt} and μ_{kt} capture fixed effects corresponding to Year × (# of Type applications) and Year × (# of Non-Type applications), respectively. (We include these year-by-count fixed effects to control for any remaining heterogeneity related to firms' total volumes of Type and Non-type patent applications per year.) Finally, the term θ_{it} captures industry-by-year fixed effects, where industry is measured with 2-digit SIC codes.

The main coefficient of interest in the above 2SLS model is δ_1 in the second-stage equation. Provided that our instruments are valid, this coefficient will capture the causal occupation-level employment effects of AI patent grants that lie within the technology category of interest.

5. Results

In this section, we first report our findings on occupational employment effects of AI innovation. We then describe the details of our empirical approach for tests of firm-level outcomes, including how we further group the seven functional types of AI patents based on their estimated employment effects. We then report our empirical findings.

5.1 Do AI technologies complement or substitute for human workers?

5.1.1 The labor-augmenting and labor-displacing effects of AI innovation

Our empirical analysis starts with the effects of different AI technologies on occupational employment growth. New technologies may increase the demand for human workers that can adapt to new production processes and operate advanced systems (Acemoglu and Restrepo, 2019; Autor and Salomons, 2018). On the other hand, some AI innovations may lead to technological unemployment due to the substitution effects (Frey and Osborne, 2017). The contrasting views could be explained by the fact that AI technologies have different functional capabilities which could affect skilled labor growth in opposite directions. For instance, the AI-based user interface implemented in the workplace can be used to collect and deliver information when interacting with employees. It may significantly increase the efficiency of existing workers without displacing any of their jobs. In addition, this technology can also increase firms' demand for other skilled workers, such as technicians for data security and system maintenance. However, some other AI technologies, such as AI-based systems that rely on machine learning and natural language processing algorithms, may disrupt existing employees by replacing repetitive tasks of human beings. They can perform reasoning and inference functions similar to those handled by data-centric workers, but in a more efficient way.

To explore the effects of AI technologies based on their functional capabilities, we implement, for each of the seven AI categories, a two-stage instrumental variables approach as described in Section 4.2. We measure firm-level AI innovations as the log-transformed number of patent grants in a particular AI category each year. To capture the direct substitutive and complementary effects of AI on jobs, we rely on the AI exposure at the technology-occupation level constructed in Section 4.1. By investigating the relation between AI innovation and firm occupational employment changes among exposed occupations, we isolate the direct AI substitution and complementarity effects from other indirect augmentation and displacement effects.

Table 5 reports the second-stage results of 2SLS regressions examining the effects of AI innovation on skilled-labor growth in occupations within firms. For ease of exposition, we report the first-stage results in Appendix C. Both the first-stage regression results and F-statistics are significant, rejecting the weak instrument hypothesis. In Panel A, the dependent variable is *Employment growth (in Headcount)*, calculated as the change in employees in an occupation scaled by the total number of employees within the firm. The *Exposed* dummy indicates whether a given occupation contains important skills related to the "type" AI as defined in Section 4.1. Each regression includes firm fixed effects, the number of "type" AI and "other" patent applications-by-year fixed effects, industry-by-year fixed effects, and a set of firm characteristics as control variables.

The significant positive coefficients on the *Log # of type k AI patents* × *Exposed* in Columns (1), (3), (5), (6), (7) of Panel A suggest that AI patents related to language, engagement, decision-making, learning, and creativity complement human workers in exposed occupations. The significant negative coefficient on *Log # of type k AI patents* × *Exposed* in Column (2) suggests that AI innovation related to perception tends to displace workers among exposed occupations. We find no significant effects of inference-related AI on overall headcount growth in exposed occupations. The results suggest that AI innovation can be either substitutive or complementary, depending on which type of AI is involved and which occupations are exposed to that type of AI.

5.1.2 Core and scope augmentation effects of AI innovation

Next, we test labor augmentation on both the intensive and extensive margins. On one hand, core augmentation for skilled labor takes place when AI innovation complements workers' existing jobs (Autor et al., 2003; Acemoglu and Autor, 2011; Acemoglu et al., 2022b) and improves the firm's production function, thus helping the firm to expand its core business. On the other hand, AI innovation could increase the scope of production, enabling firms to explore new products and businesses and create new jobs (Cockburn et al., 2019; Babina et al., 2024).

To explore whether AI increases labor demand by increasing the scope of labor inputs or enabling firms to scale up the existing labor force, we test *Employment growth (with new skills)*, defined as the change in number of employees with new skills in an occupation, scaled by the total number of employees within the firm. We define an employee as having new skills if his or her skill set contains unique skills that the occupational median skill set does not have. In Panel B, the significant positive coefficients on the *Log # of type k AI patents* × *Exposed* in Columns (3), (6), (7) suggest that AI patents related to engagement, learning, and creativity complement human workers in exposed occupations by increasing the scope of worker skills. The significant negative coefficients on *Log # of type k AI patents* × *Exposed* in Column (2) suggest that perception-based AI innovation substitutes for workers in exposed occupations by reducing the skills required for the job. We find no significant effects of language AI and decision-making AI on labor growth with new skills, suggesting that language AI and decision-making AI mainly complement human work with existing skills. Overall, the results are consistent with the recent literature on both skill-technology complementarity and skill-technology substitution. Also, the finding that different types of AI innovation can have very heterogeneous effects on occupation-level employment helps to provide a more nuanced perspective in the ongoing debate over whether AI technology ultimately drives labor augmentation or labor displacement.

5.2 AI innovation and firm efficiency

Our previous findings suggest that AI innovation significantly impacts firm-level employment dynamics by both complementing and substituting human workers in exposed occupations. At the firm level, labor augmentation and displacement may involve different cost-benefit tradeoffs, leading to implications for firm efficiency. To gain more insights, this section further explores the direct consequences of AI innovations that augment and displace human workers. We expect that labor augmentation and displacement may have effects on firm efficiency through different channels. For instance, when adopting AI technologies increases labor growth, firm efficiency can be improved due to productivity increases. At the same time, in the case of labor growth decreases resulting from AI technologies, firms might also benefit due to a reduction in labor costs. To shed light on the importance of these two potential channels, we explore the heterogeneity of AI technologies by re-classifying AI patents into larger groupings based on their predicted effects on firms' employment outcomes.

5.2.1 Labeling patents based on firm employment effects.

To study the effects of AI innovation on firm efficiency, we re-classify all AI patents into *labor augmenting* AI patents and *labor displacing* AI patents based on predicted firm employment growth in headcounts. Specifically, we identify whether patents in each AI category filed by a firm in year t are predicted to increase or decrease the overall occupational employment growth in the following year. Based on the regression results in Panel A of Table 5, we define all AI patents except for perception-based AI and inference-based AI as labor-augmenting patents. All perception-based AI patents are defined as labor-displacing patents. Further, we classify engagement-, learning-, and creativity-based AI as *scope-augmenting* patents based on the results in Panel B of Table 5. The remaining categories of AI patents—language and decision-making—bring about increased headcount employment but do not lead to increased scope. Hence, we define the language and decision-making AI patents as *core-augmenting* AI patents. Based on this grouping, we construct the measure *Log # augmenting (displacing)* as the log-transformed total number of patent grants that are in any of the AI technologies predicted to increase (reduce) employment within exposed occupations. Similarly*, the Log # scope (core)-augmenting* is the log-transformed total number of patent grants that are (are not) predicted to augment labor with new skills. We also re-construct instruments for each type of patent using a similar approach to the one described in Section 4.2.

5.2.2 The effects of AI patents on firm productivity

Using the labeled patents and reconstructed instruments, we repeat the IV-2SLS analysis in firm-level tests and study firms' productivity changes caused by AI innovation. We use total factor productivity (TFP) as a measure of firm efficiency. We obtain a firm-level TFP measure from the residual estimates in a firm-year level linear production function model as used in prior studies (e.g., Nadiri, 1970; Craig and Harris, 1973) with capital and labor inputs, controlling for firm fixed effects. Firm-level production inputs from capital and labor are separately proxied for by cost of goods sold (from Compustat) and firms' total employment (from Revelio Labs). The dependent variable in the residual estimation regression is the firm's production output as measured by net sales (from Compustat).

Table 6 reports the second-stage results of 2SLS regressions examining the effects of AI innovation on firm value. The first-stage regression results are reported in Appendix C. The dependent variable, *Total Factor Productivity* $_{t,\ t+1}$, is the log-transformed change in the yearly TFP. The residual estimates of TFP are scaled in thousands before the log transformation. The independent variables are "type" and "other" patent grants in year t. Columns (1)-(4) report the results from OLS regressions, and Columns (5)-(8) report second-stage 2SLS results. From the OLS regression results, we find that augmenting AI patents, including both scope- and core-augmenting AI patents, have significant positive correlations with firms' subsequent productivity changes. Furthermore, augmenting AI patents—including both scope- and core-augmenting AI patents—have significant positive effects once potential endogeneity is accounted for via the 2SLS regressions. Overall, the results in Table 8 confirm the view that AI innovation, either through core augmentation or scope augmentation, can increase overall firm productivity. However, there is no evidence to suggest that labor-displacing AI innovation brings about increases in total factor productivity.

5.2.3 The effects of AI patents on cost savings

Next, to further investigate the effects of labor-displacing AI innovation, we examine changes in firms' cost savings. Table 7 reports the results of regressions relating displacing AI to firm expenditures. We use two measures of costs related to firms' production and operations: the percentage growth rates from t to t+1 in total operating expenses and in SG&A expenditures. The independent variables in the regressions are "type" and "other" patent grants in year t. Columns (1) and (3) report the results from OLS regressions, and Columns (2) and (4) report 2SLS second-stage regression results. As seen in the table, the OLS regression results do not show significant coefficients for the log number of displacing AI patents. However, Columns (2) and (4) reveal that, once endogeneity is accounted for via our instrumental variables estimation, significant negative coefficients emerge on *Log # Displacing*. Taken together with the earlier results on displacing AI from Table 6, the results in Table 7 suggest that the benefits to firms from labor-displacing AI largely stem from cost reductions rather than productivity gains.

5.3 AI innovation and firm value

5.3.1 The effects of AI patents on firm value

Next, we examine firm value changes resulting from AI innovation. We expect AI innovation to bring value increases to firms through either productivity enhancements and, in the case of displacing innovations, cost reductions. Table 8 reports the results of 2SLS regression examining the effects of AI innovation on firm value. The dependent variable is the net change in Tobin`s Q from year t to t+1. The independent variables are "type" and "other" patent grants in year t. Columns (1)-(4) report the results from OLS regressions and Columns (5)-(8) report the second-stage 2SLS regression results. (First-stage regression results are reported in Appendix C.) The significant results in all columns show that augmenting AI (through core and scope augmentation) and displacing AI alike have significant positive effects on firm value.

5.3.2 Heterogeneity of AI innovation effects on firm value

In this section, we explore the effects of AI innovation conditional on different labor market frictions that a firm may face. Specifically, we study (1) how the value effects of augmenting AI depend on a firm's hiring costs and (2) how the value effects of displacing AI are shaped by a firm's termination costs.

As discussed in the previous section, AI innovation can increase firm value through labor augmentation because new technology improves the firm's growth by augmenting employees (with or without new skills). The potential value effects of AI-induced labor augmentation can vary with the environment of the external labor market, especially when AI increases firms' demand for new skills. We expect that, in general, augmenting AI innovations are more value-creating when firms have lower costs associated with hiring activities. We rely on three measures of hiring costs to conduct subsample analyses. First, we use the average employee turnover rate among local firms within the industry (excluding the focal firm) as a proxy for local job candidate mobility. Firms located in areas with high labor mobility may find it easier to hire new employees. Second, we use the state-level enforceability index for non-compete clauses (NCC) as another measure of external hiring frictions. Prior studies have documented that the increase in non-compete enforceability can restrict labor mobility and induce firms' labor costs, especially for knowledge-intensive occupations (Starr et al. 2021; Jeffers, 2024). On one hand, non-compete enforceability captures the ease with which firms can hire employees from rivals. On the other hand, when non-compete agreements are enforceable, both employee retention costs and termination costs are lower. Third, we construct a measure of within-occupation transferability as a proxy for the costs of hiring employees into occupations from the same occupations. When firm demand for existing labor increases, high within-occupation transferability can reduce hiring costs.

The results for the tests involving labor market frictions are reported in Table 9. The dependent variable is the net change in Tobin`s Q from year t to t+1. The independent variables are "type" and "other" patent grants in year t. Panel A reports regression results for subsamples of firms headquartered in high (above the sample median) employee turnover states and low employee turnover states. *Employee Turnover Rate* is defined as the state-level average job separations scaled by the firm total employees excluding the focal firm. In Column (1), we find significant positive effects of overall labor-augmenting AI patents on firm value, compared to insignificant effects in Column (2). As reported in Column (3), we find significant positive effects of scope-augmenting AI patents on firm value, compared to the results in Column (4). Specifically, augmenting (scope-augmenting) AI innovation increases firm value when firms are in states with high employee turnover rates (mobility). We find no significant effects in Columns (5)-(6) for core-augmenting AI. Thus, when comparing the results for scope- and core-augmenting AI patents, it appears that firms' value gains are more dependent on external labor market frictions for the former. This can be explained by the fact that employees with new skills are typically hired from outside of the firm, and scope augmentation is more sensitive to the costs of searching for, recruiting, and hiring more specialized talent. Panel B reports the results in subsamples of firms with headquarters located in high (above the sample median) non-compete enforceability and low NCC enforceability states. *NCC enforceability* is the firm's headquarters-state enforceability index for non-compete clauses (NCC), drawn from Garmaise (2011). We find that both scope-augmenting and core-augmenting AI patents have significant positive effects on value when firms are located in high NCC enforceability states. Panel C reports the results in subsamples of firms in high (above the sample median) within-occupation transferability industries and low within-occupation transferability industries. The results suggest that the benefits of employee retention (due to lower risks of losing trade secrets) in high enforceability states dominate the increased hiring costs (due to the difficulty of hiring from rivals).[25] *Within-occupation transferability* is the industry-level (2-digit SIC) weighted average within-occupation transferability among occupations exposed to AI. For each occupation, we calculate the within-occupation transferability as the percentage of job inflow transitions within 3-digit SOC occupations scaled by total inflow transitions. We find significant positive value effects of scope-augmentation for firms with low within-occupation transferability. In fact, these firms are likely to have high cross-occupation transferability, which benefits firms more when new skills are needed. In Column (5), the positive coefficients on *Log # Core augmenting* suggest that core augmentation creates more value when within-occupation transferability is high. Overall, the results in Table 9 are consistent with the prediction that firms benefit more from augmenting AI when they have fewer hiring costs. In particular, whereas external labor mobility relaxes the hiring constraints for scope-augmenting AI that, by definition, requires new skills, within-occupation mobility is more beneficial for firms that obtain core-augmenting AI innovations—presumably because core augmentation does not often involve cross-occupation worker transitions.

Next, to investigate the heterogeneous effects of displacing AI under different termination costs, we focus on two aspects of labor protection. First, we use state-level unemployment insurance as a measure of labor adjustment costs. High unemployment insurance benefits are a form of protection against unemployment risk for workers, but such benefits can increase the costs that firms must bear in the event of firings and layoffs. In addition, employees in firms with less generous unemployment insurance benefits bear high unemployment risks and demand a wage premium as compensation (Topel, 1984a). We expect that firms located in states with low unemployment insurance benefits are subject to lower termination costs. Second, terminating workers may be harmful to the firm due to the loss of firm-specific information or trade secrets, given that a key channel of inter-firm information diffusion is through employee turnover. This can be alleviated if labor mobility is limited. Accordingly, in studying value effects of labor-

[25] This is especially true when scope-augmenting AI increases the demand for new skills, which is less dependent on labor supply from direct competitors.

displacing AI, we again use NCC enforceability as a restriction on labor mobility. Overall, we expect displacing AI innovation to create more value when firms are located in states with high restrictions on labor mobility.

We report the subsample test results in Table 10. Columns (1) and (2) report the results in subsamples of firms with headquarters located in states with high versus low unemployment insurance benefits. The state-level unemployment insurance data are obtained from the U.S. Department of Labor.[26] We measure unemployment insurance benefits as the product of average weekly benefits and the recipiency rates from the regular UI programs. In Columns (3) and (4), we consider subsamples of firms with headquarters located in high versus low NCC enforceability states. We find significant positive coefficients on *Log # Displacing* in Columns (2) and (3), but not in Columns (1) and (4). This shows that the effects of displacing AI innovation on firm value are affected by firms' exposure to termination costs and frictions in the external labor market. Specifically, firms with lower costs of worker termination (e.g., less unemployment protection and restricted labor mobility) can gain more value from adopting labor-displacing AI technologies.

Lastly, we study how the distribution of worker skills in the external and internal labor markets influences the positive relation between augmenting AI and firm value. We focus on two dimensions along which firms' skill accumulation process might be affected. First, we consider the skill-based proximity among occupations. We construct an occupation-pair-level skill distance by counting the number of unique skills (weighted by skill importance from the O*NET database) in the occupation relative to all other pair occupations. Then we aggregate the skill distance to the industry and state level by taking the average (weighted by occupational employment share) skill distance across occupations exposed to "type" AI innovations. An industry or state with high skill distance is taken to have low skill proximity. Intuitively, proximate skills are more transferable across occupations, thus lowering search/recruiting costs and helping to facilitate augmentation. Second, we focus on the supply of potential skills that can help meet augmenting firms' labor demand from both internal and external markets. We calculate the occupation-level skill supply by counting the number of important skills (O*NET importance score greater than 4) required for an occupation. Then we aggregate the skill supply to the industry and firm level by taking the average (weighted by occupational employment shares) across occupations that are exposed to a given type

---

[26] The data source is from https://oui.doleta.gov/unemploy/data.asp

of AI innovation. We expect that augmenting AI innovation benefits firms more when the labor markets offer high skill proximity and can provide a sufficient supply of skills.

Table 11 reports the results of subsample tests based on skill proximity and skill supply. In Panel A, Columns (1), (3), and (5) report results for the subsample of firms that are either located in a state with high skill proximity or operate in a (2-digit SIC) industry with high skill proximity. Columns (2), (4), and (6) report the results for the subsample of firms that are located in a state with low skill proximity and also operate in a (2-digit SIC) industry with low skill proximity. Consistent with expectations, the positive coefficients on *Log # (Scope-/Core-) Augmenting in* Columns (1), (3), and (5) show that suggest that skill proximity has a positive moderating effects for the value impact of augmenting AI innovation. In Panel B, Columns (1), (3), and (5) report the results for the subsample of firms that have either a high level of internal skill supply or operate within a (2-digit SIC) industry that has a high level of skill supply. Columns (2), (4), and (6) report the results for the subsample of firms that have a low internal skill supply and also operate in a (2-digit SIC) industry with a low skill supply. The positive coefficients in Columns (1), (3), and (5) suggest that a higher supply of related and required skills can enhance the positive value effects of labor-augmenting AI innovation. Overall, the evidence in Table 11 highlights the important role of the labor-market distribution of skills in amplifying the positive value effects of AI augmentation.

6. Conclusion

The rapid pace of innovation in the field of AI has led to growing interest among academics, practitioners, and policymakers in understanding the possible benefits and consequences of new technologies in artificial intelligence that can process vast amounts of information, perform cognitive tasks such as optimization, learning, reasoning, and prediction, generate novel content, and engage with human users. While increasingly powerful AI technologies offer the promise of improving firms' productivity within virtually every industrial sector, questions have arisen about the effects that AI will have on labor markets. Does AI tend to substitute and complement human workers, thus creating new roles, new jobs, and demand for new skills? Or, will AI be able to match and surpass humans in cognitive workplace tasks, leading to skill obsolescence, labor displacement, and technological unemployment? These questions have been difficult to answer due to the lack of large-scale data on the various types of AI innovation and the impact that these innovations have on the movements of individual workers.

In this paper, we use over five million patent filing texts from 2007-2023 and microdata on workforce dynamics to study the labor-market and firm-efficiency impacts of AI innovation. To implement our study, we propose a new categorization of AI based on seven functional areas that correspond to key aspects of cognitive domains in human intelligence: language, perception, inference, engagement, decision-making, learning, and creativity. Using state-of-the-art techniques in Large Language Models (LLMs) and Generative AI to identify patents in these functional categories, we study the causal impact of each innovation type on occupational employment growth and firm outcomes. By unpacking the AI technologies and occupational AI exposure, we document evidence of both augmentation and displacement: AI innovations related to engagement, creativity, language, learning, and decision-making significantly increase labor growth, while AI innovations related to perception significantly reduce it. Further, we find that engagement AI, learning AI, and creativity AI not only augment labor in terms of overall headcount, but also increase the net hiring of workers with new skills, thus facilitating scope augmentation.

While augmenting AI innovations generally translate into higher productivity and firm valuations, their exact value impact depends heavily on the labor market circumstances surrounding the innovating firm. We find, for example, that the positive value impact of labor-augmenting AI innovations can be enhanced when firms face lower external hiring costs as measured by high labor mobility, high within-occupation transferability, high skill proximity, and a sufficient supply of new skills. In addition, labor displacing AI can lead to value increases due to cost-saving benefits, especially when firms face fewer external labor market costs from worker terminations (e.g., weak unemployment protection and restrictions on labor mobility). Overall, our findings are consistent with the view that many AI innovations have favorable effects for both innovating firms and the affected workforce, and these positive effects may be greatly facilitated when the costs of labor adjustment are low.

References

Abis, S., and Veldkamp, L. 2024. The changing economics of knowledge production. *The Review of Financial Studies 37*:89-118.

Acemoglu, D. 2002. Technical change, inequality, and the labor market. *Journal of Economic Literature*, 40:7-72.

Acemoglu, D., and Autor, D. 2011. Skills, tasks and technologies: Implications for employment and earnings. *Handbook of Labor Economics* 4:1043-1171.

Acemoglu, D., and Restrepo, P. 2018. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review 108*:1488-1542.

Acemoglu, D., Restrepo, P. 2019. Automation and new tasks: how technology displaces and reinstates labor. *J. Econ. Perspect* 33:3–30.

Acemoglu, D., and Restrepo, P. 2020. Robots and jobs: Evidence from U.S. labor markets. *Journal of Political Economy 128*:2188-2244.

Acemoglu, D., and Shimer, R. 1999. Holdups and efficiency with search frictions. *International Economic Review* 40*:*827-849.

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. 2022b. Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics* 40: S293-S340.

Aghion, P., Jones, B. F., and Jones, C. I. 2017. Artificial intelligence and economic growth (Vol. 23928). Cambridge, MA: *National Bureau of Economic Research.*

Agrawal, A., Gans, J., and Goldfarb, A. 2019a. *The economics of artificial intelligence: an agenda.* Chicago, IL: University of Chicago Press.

Agrawal, A., Gans, J. S., and Goldfarb, A. 2019b. Artificial intelligence: the ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives* 33:31-50.

Alderucci, D., Branstetter, L., Hovy, E., Runge, A., and Zolas, N. 2020. Quantifying the impact of AI on productivity and labor demand: Evidence from US census microdata. *In Allied social science associations—ASSA 2020 annual meeting*.

Austin, D. H. 1993. An event-study approach to measuring innovative output: The case of biotechnology. *The American Economic Review* 83:253-258.

Autor, D. H., Levy, F., and Murnane, R. J. 2003. The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics* 118*:*1279-1333.

Autor, D., and Salomons, A. 2018. Is automation labor-displacing? Productivity growth, employment, and the labor share. NBER Working Paper No. w24871.

Babina, T., Fedyk, A., He, A., and Hodson, J. 2023. *Firm investments in artificial intelligence technologies and changes in workforce composition*. National Bureau of Economic Research Vol. 31325.

Babina, T., Fedyk, A., He, A., and Hodson, J. 2024. Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics* 151:103745.

Badet, J. 2021. AI, automation and new jobs. *Open Journal of Business and Management* 9:2452-2463.

Bates, T. W., Du, F., and Wang, J. J. 2021. Workplace automation and corporate financial policy. Working Paper, Arizona State University.

Bennett, B., Stulz, R., and Wang, Z. 2020. Does the stock market make firms more productive? *Journal of Financial Economics*, 136(2), 281-306.

Berman, E., Bound, J., and Machin, S. 1998. Implications of skill-biased technological change: international evidence. *The Quarterly Journal of Economics 113:*1245-1279.

Bessen, J. 2018. AI and jobs: The role of demand. NBER Working Paper No. w24235.

Brynjolfsson, E., and McAfee, A. 2017. The Business of Artificial Intelligence, *Harvard Business Review*, July 18, 2017.

Brynjolfsson, E., Mitchell, T., and Rock, D. 2018. What can machines learn and what does it mean for occupations and the economy? *AEA Papers and Proceedings* 108:43-47.

Caliskan, A., and Lum, K. 2024. Effective AI regulation requires understanding General-Purpose AI. *Brookings Institution* January 2024.

Cao, S., Jiang, W., Wang, J. L., and Yang, B. 2024. From man vs. machine to man+ machine: The art and AI of stock analyses. *Journal of Financial Economics Forthcoming.*

Chava, S., Danis, A., and Hsu, A. 2020. The economic impact of right-to-work laws: Evidence from collective bargaining agreements and corporate policies. *Journal of Financial Economics* 137:451-469.

Chen, M. A., Wu, Q., and Yang, B. 2019. How valuable is FinTech innovation? *The Review of Financial Studies* 32:2062-2106.

Cockburn, I., Henderson, R. and Stern, S. 2019. The Impact of Artificial Intelligence on Innovation: An Exploratory Analysis. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press 115-148.

Craig, C. E., & Harris, R. C. 1973. Total productivity measurement at the firm level. *Sloan Management Review* (pre-1986), 14(3), 13.

Devos, E., and Rahman, S. 2018. Labor unemployment insurance and firm cash holdings. *Journal of Corporate Finance*, *49*, 15-31.

Ellul, A., Erel, I., and Rajan, U. 2020. The COVID-19 pandemic crisis and corporate finance. *The Review of Corporate Finance Studies* 9:421-429.

Erickson, T., and Whited, T. M. 2012. Treating measurement error in Tobin's q. *The Review of Financial Studies* 25:1286-1329.

Eisfeldt, A. L., Schubert, G., Zhang, M. B., and Taska, B. 2024. The labor impact of generative AI on firm values. Working Paper, UCLA Anderson School of Management.

Ernst, E., Merola, R., and Samaan, D. 2019. Economics of artificial intelligence: Implications for the future of work. *IZA Journal of Labor Policy* 9:1.

Ersahin, N. 2020. Creditor rights, technology adoption, and productivity: Plant-level evidence. *The Review of Financial Studies 33*:5784-5820.

Farre-Mensa, J., Hegde, D., and Ljungqvist, A. 2020. What is a patent worth? Evidence from the US patent "lottery". *The Journal of Finance* 75:639-682.

Fazzari, S., Hubbard, R. G., and Petersen, B. 1988. Investment, financing decisions, and tax policy. *The American Economic Review* 78:200-205.

Frey, C. B., and Osborne, M. A. 2017. The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change* 114:254-280.

Furman, J., and Seamans, R. 2019. AI and the Economy. *Innovation Policy and The Economy* 19:161-191.

Garmaise, M. J. 2011. Ties that truly bind: Noncompetition agreements, executive compensation, and firm investment. *The Journal of Law, Economics, & Organization*, 27:376-425.

Gaule, P. 2018. Patents and the success of venture-capital backed startups: Using examiner assignment to estimate *causal* effects. *The Journal of Industrial Economics* 66:350-376.

Ghaly, M., Anh Dang, V., and Stathopoulos, K. 2017. Cash holdings and labor heterogeneity: The role of skilled labor. *The Review of Financial Studies* 30:3636-3668.

Giczy, A. V., Pairolero, N. A., and Toole, A. A. 2022. Identifying artificial intelligence (AI) invention: A novel AI patent dataset. *The Journal of Technology Transfer* 47:476-505.

Grennan, J., and Michaely, R. 2020. Artificial intelligence and high-skilled work: Evidence from analysts. *Swiss Finance Institute Research Paper* 20-84.

Griliches, Z. 1969. Capital-skill complementarity. *The Review of Economics and Statistics* 465-468.

Hall, B. H., Jaffe, A. B., and Trajtenberg, M. 2001. The NBER patent citation data file: Lessons, insights and methodological tools. NBER Working Paper No. w8498.

Hanley, K. W., and Hoberg, G. 2010. The information content of IPO prospectuses. *The Review of Financial Studies* 23:2821-2864.

Henry, E., and Leone, A. J. 2016. Measuring qualitative information in capital markets research: Comparison of alternative methodologies to measure disclosure tone. *The Accounting Review* 91:153-178.

Hoberg, G., and Phillips, G. 2016. Text-based network industries and endogenous product differentiation. *Journal of Political Economy* 124:1423-1465.

Hsu, P. H., Tian, X., and Xu, Y. 2014. Financial development and innovation: Cross-country evidence. *Journal of Financial Economics* 112:116-135.

Huang, Q., Shen, Y., Sun, Y., and Zhang, Q. 2023. The layoff generation: How generative AI will reshape employment and labor markets. Working Paper, The University of Illinois at Urbana-Champaign.

Hussey, R. 2003. Labor turnover and the dynamics of labor productivity. Working Paper, Georgetown University.

Jeffers, J. S. 2024. The impact of restricting labor mobility on corporate investment and entrepreneurship. *The Review of Financial Studies* 37:1-44.

Jegadeesh, N., and Wu, D. 2013. Word power: A new approach for content analysis. *Journal of Financial Economics* 110:712-729.

Jha, M., Qian, J., Weber, M., and Yang, B. 2024. ChatGPT and corporate policies. NBER Working Paper No. w32161.

Kogan, L., Papanikolaou, D., Seru, A., and Stoffman, N. 2017. Technological innovation, resource allocation, and growth. *The Quarterly Journal of Economics* 132:665-712.

Kogan, L., Papanikolaou, D., Schmidt, L. D., and Seegmiller, B. 2023. Technology and labor displacement: Evidence from linking patents with worker-level data. NBER Working Paper No. w31846.

Krusell, P., Ohanian, L. E., Ríos-Rull, J. V., and Violante, G. L. 2000. Capital-skill complementarity and inequality: A macroeconomic analysis. *Econometrica* 68:1029-1053.

Kumar, I. 2023. Impact of Robo-advisors on the Labor Market for Financial Advisors. Working Paper, Emory University.

Lane, M., and Saint-Martin, A. (2021). The impact of Artificial Intelligence on the labour market: What do we know so far? *OECD Social, Employment and Migration Working Papers* No. 256

LeCun, Y. 2022. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. *Open Review* 62.

Lemley, M. A., and Sampat, B. 2012. Examiner characteristics and patent office outcomes. *Review of Economics and Statistics*, 94:817-827.

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems* 33:9459-9474.

Li, K., Mai, F., Shen, R., and Yan, X. 2021. Measuring corporate culture using machine learning. *The Review of Financial Studies* 34:3265-3315.

Li, K., Mai, F., Shen, R., Yang, C., and Zhang, T. 2024. Dissecting corporate culture using generative AI–Insights from analyst reports. Working Paper, University of British Columbia.

Loughran, T., and McDonald, B. 2011. When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *The Journal of Finance* 66:35-65.

Matsa, D. A. 2010. Capital structure as a strategic variable: Evidence from collective bargaining. *The Journal of Finance* 65*:*1197-1232.

McCarthy, J., Minsky, M. L., Rochester, N., and Shannon, C. E. 2006. A proposal for the Dartmouth summer research project on artificial intelligence, August 31, 1955. *AI magazine 27*:12-12.

Melero, E., Palomeras, N., and Wehrheim, D. 2020. The effect of patent protection on inventor mobility. *Management Science* 66:5485-5504.

Nadiri, M. I. 1970. Some approaches to the theory and measurement of total factor productivity: a survey. *Journal of Economic Literature*, 8(4), 1137-1177.

Pakes, A. 1985. On patents, R&D, and the stock market rate of return. *Journal of Political Economy* 93:390-409.

Qiu, J., Wan, C., and Wang, Y. 2021. Labor-capital substitution and capital structure: Evidence from automation. Working Paper, McMaster University.

Rossi, A. G., and Utkus, S. P. 2021. Who benefits from Robo-advising? Evidence from machine learning. Working Paper, Georgetown University.

Routledge, B. R., Sacchetto, S., and Smith, N. A. 2013. Predicting merger targets and acquirers from text. Working Paper, Carnegie Mellon University.

Sampat, B., and Williams, H. L. 2019. How do patents affect follow-on innovation? Evidence from the human genome. *American Economic Review* 109:203-236.

Seamans, R., and Raj, M. 2018. AI, labor, productivity and the need for firm-level data. NBER Working Paper No. w24239.

Starr, E. P., Prescott, J. J., and Bishara, N. D. 2021. Noncompete agreements in the US labor force. *The Journal of Law and Economics* 64:53-84.

Tetlock, P. C., Saar-Tsechansky, M., and Macskassy, S. 2008. More than words: Quantifying language to measure firms' fundamentals. *The Journal of Finance* 63:1437-1467.

Topel, R. H. 1984a. Equilibrium earnings, turnover, and unemployment: New evidence. *Journal of Labor Economics* 2:500-522.

Turing, A. M. 1950. Computing machinery and intelligence. *Mind* 49:433-460.

Van Doornik, B., Fazio, D., Schoenherr, D., and Skrastins, J. 2022. Unemployment insurance as a subsidy to risky firms. *The Review of Financial Studies* 35:5535-5595.

Webb, M. 2019. The impact of artificial intelligence on the labor market. Working Paper, Stanford University.

Yang, C. H. 2022. How artificial intelligence technology affects productivity and employment: firm-level evidence from Taiwan. *Research Policy* 51:104536.

Yang, B., and Yuan, T. 2022. Trademark and IPO underpricing. *Financial Management* 51:271-296.

Appendix A. Examples of AI Innovation Types Identified by LLM Queries

Language

Query: *Is the following invention directly related to artificial intelligence that has language ability? Respond with YES or NO and give an explanation.*

Title: AUTOMATIC CLASSIFICATION OF DRILLING REPORTS WITH DEEP NATURAL LANGUAGE PROCESSING
Patent Application Number: 20210165963
Patent Applicant: Landmark Graphics Corporation

Abstract: Systems, methods, and computer-readable media for automatic classification of drilling reports with deep natural language processing. A method may involve obtaining drilling reports associated with respective well drilling or operation activities, and based on the drilling reports, generating a plurality of word vectors, wherein each word vector from the plurality of word vectors represents a respective word in the drilling reports. The method can further involve partitioning sentences in the drilling reports into respective words and, for each sentence, identifying respective word vectors from the plurality of word vectors, the respective word vectors corresponding to the respective words associated with the sentence. The method can involve classifying via a neural network, the sentences in a drilling report into at least one of respective events, respective symptoms, respective actions, and respective results. The method can also classify sentences according to any set of labels of interest.

Explanation from LLM:
YES. This invention is directly related to artificial intelligence with language abilities because it involves the use of deep natural language processing (NLP) techniques to automatically classify content within drilling reports. Specifically, it uses word vectors and neural networks to process and categorize sentences from drilling reports into different categories such as events, symptoms, actions, and results. This kind of task requires the AI system to understand and interpret human language, which is a key aspect of AI with language abilities.

Title: FURTHER IMPROVEMENTS IN RECOMMENDATION SYSTEMS
Patent Application Number: 20190026812
Patent Applicant: 4-TELL, Inc.

Abstract: This invention deals with improving recommendation systems. The first embodiment combines rules and recommendations to create automated and intelligent business rules for recommendations. The second embodiment improves recommendations by combining the results of driver products and influencer products, where influencer products only influence the recommendations of the driver products. Influencer products can be related to a specific user. The third embodiment improves recommendations for new items by relating them to original items, such that the sales for the original item is used in the new item when calculating recommendations. The new items may replace the original item, or be a similar item and exist alongside the original item.

Explanation from LLM:
NO. While this invention enhances recommendation systems which can be part of AI applications, it does not specifically mention or require a system with language abilities. The improvements described focus on how recommendations are generated and refined through various strategies (such as combining rules, considering driver and influencer products, and relating new items to original items), but these processes do not inherently involve natural language processing or generation, which would be necessary for a system to be considered directly related to AI with language abilities.

Perception

Query: *Is the following invention directly related to artificial intelligence that has perceptual ability? Respond with YES or NO and give an explanation.*

<u>Title: USER INTERACTIONS IN MOBILE DAMAGE ASSESSMENT AND CLAIMS PROCESSING</u>
Patent Application Number: 20220058743
Patent Applicant: Allstate Insurance Company

Abstract: Systems and methods provide for an automated system for analyzing damage and processing claims associated with an insured item, such as a vehicle. An enhanced claims processing server may analyze damage associated with the insured item using photos/video transmitted to the server from a user device (e.g., a mobile device). The mobile device may receive feedback from the server regarding the acceptability of submitted photos/video, and if the server determines that any of the submitted photos/video is unacceptable, the mobile device may capture additional photos/video until all of the data are deemed acceptable. To aid in damage analysis, the server may also interface with various internal and external databases storing reference images of undamaged items and cost estimate information for repairing previously analyzed damages to similar items. Further still, the server may generate a payment for compensating a claimant for repair of the insured item.

Explanation from LLM:
YES. The described invention involves an automated system that analyzes damage to insured items, such as vehicles, using photos or videos provided by users. This process requires the

system to interpret visual data, which can be considered a form of perceptual ability within the realm of artificial intelligence. Specifically, the system must analyze images and videos to assess damage, compare them with reference images of undamaged items, and possibly even estimate repair costs. These tasks involve pattern recognition, image analysis, and potentially machine learning techniques, which are key components of AI systems with perceptual abilities. Therefore, this invention is directly related to artificial intelligence that has perceptual ability.

Title: USER INTERACTIONS IN MOBILE DAMAGE ASSESSMENT AND CLAIMS PROCESSING
Patent Application Number: 20230007284
Patent Applicant: Google LLC

Abstract: Ultra light models and decision fusion for increasing the speed of intra-prediction are described. Using a machine-learning (ML) model, an ML intra-prediction mode is obtained. A most-probable intra-prediction mode is obtained from amongst available intra-prediction modes for encoding the current block. As an encoding intra-prediction mode, one of the ML intra-prediction mode or the most-probable intra-prediction mode is selected, and the encoding intra-prediction mode is encoded in a compressed bitstream. A current block is encoded using the encoding intra-prediction mode. Selection of the encoding intra-prediction mode is based on relative reliabilities of the ML intra-prediction mode and the most-probable intra-prediction mode.

Explanation from LLM:
NO. The described invention involves the use of a machine learning (ML) model for obtaining an intra-prediction mode in video compression, but it does not inherently involve perceptual abilities. Perceptual abilities typically refer to the capability of recognizing and interpreting sensory information such as images, sounds, or other stimuli, which often requires more complex cognitive processing. In this case, the ML model is used to improve the efficiency of the encoding process by selecting the best intra-prediction mode among available options, rather than perceiving and understanding visual content in a way that mimics human perception. This application of ML is focused on optimizing a specific technical aspect of video compression rather than developing perceptual capabilities.

Engagement

Query: *Is the following invention directly related to artificial intelligence that has the ability to engage with users? Respond with YES or NO and give an explanation.*

Title: USING AUGMENTED REALITY TO DEPICT PRE-QUALIFIED PURCHASES
Patent Application Number: 20210110460
Patent Applicant: Wells Fargo Bank, N.A.

Abstract: A list of products that a user is pre-qualified to purchase is received at an augmented reality device from an electronic computing device. One or more virtual objects are received from the electronic computing device. Each of the virtual objects is associated with one of the products that the user is pre-qualified to purchase. An image is received of a view obtained via the augmented reality device. One or more positions are identified in the view in which to display one or more of the virtual objects. The virtual objects are displayed in the identified one or more positions on the image of the view. A selection is received of a virtual object of the one or more virtual objects. The virtual object is associated with a product that the user is pre-qualified to purchase. A purchase of the product is initiated.

Explanation from LLM:
YES. This invention involves the use of augmented reality (AR) technology to display virtual objects representing products that a user is pre-qualified to purchase within their real-world view captured by an AR device. This interaction process inherently requires some form of AI to handle the pre-qualification of products for the user, to identify suitable positions within the captured view to place the virtual objects, and to manage the user's interaction with these objects, leading to a potential purchase. The system likely uses machine learning algorithms to understand user preferences, predict which products might interest them, and to improve the user experience through continuous learning and adaptation. Therefore, it is directly related to AI as it engages with users in a personalized manner to facilitate commerce through augmented reality interactions.

Title: DYNAMICALLY RECONFIGURABLE INSURANCE PRODUCT
Patent Application Number: 20210166322
Patent Applicant: State Farm Mutual Automobile Insurance Company

Abstract: A dynamically reconfigurable insurance product, system, and computer-implemented method may, with customer permission or consent, collect customer data; analyze the customer data to determine life events, and customer location and activities; and dynamically adjust the insurance product in real-time or substantially real-time. The dynamically reconfigurable insurance product may include several types of insurance, such as auto, home, life, personal articles, etc. From the data collected, risk levels associated with the insured, their family, and/or personal belongings may be adjusted. Based upon the risk levels determined, different types of the insurance within the insurance product may be updated, or new types of insurance may be added to the insurance product. For instance, based upon a marriage or birth of a child, life insurance coverage may be added or increased. The customer may then be notified of the changes, or proposed changes, and approve or reject the changes to the insurance product.

Explanation from LLM:

NO. While this invention involves sophisticated data collection and analysis techniques, it is primarily focused on automating and personalizing insurance products rather than engaging directly with users in a manner typical of AI applications like chatbots or virtual assistants. The system described does interact with customers by notifying them of changes and seeking approval for adjustments, but its core function is not centered around direct user interaction or engagement through conversational interfaces or similar AI-driven methods. Instead, it uses AI for analyzing data and making decisions about insurance products.

Inference

Query: *Is the following invention directly related to artificial intelligence that has inference ability? Respond with YES or NO and give an explanation.*

Title: AUTOMATED SELF-SERVICE USER SUPPORT BASED ON ONTOLOGY ANALYSIS
Patent Application Number: 20180218067
Patent Applicant: International Business Machines Corporation

Abstract: A system for implementing a method that provides information to a user in response to a received user query. A natural language analysis generates substrings relevant to the user query. An ontology analysis outputs: terms of an ontology matching the relevant generated substrings; and relationships between the terms. A query analysis analyzes the user query regarding the outputted terms and relationships, including ascertaining whether the user query is more suitable for service than for an information search. If it is so ascertained, then service actions for the user to perform are identified to the user. If it is not so ascertained, then: the user query is refined based on the outputted terms and relationships; a search query is generated based on the refined user query, a search is initiated based on the search query, and results of the search are provided to the user.

Explanation from LLM:
YES. The described invention involves a system that uses natural language processing (NLP) and ontology analysis to understand user queries and provide appropriate responses. This process includes analyzing the user's input, understanding its context through ontological relationships, and determining whether the query is better suited for providing services or performing an information search. The system's capability to analyze and infer meaning from natural language inputs, refine those inputs, and generate actions or searches based on the refined queries indicates a sophisticated use of AI technologies that include machine learning and reasoning capabilities. Therefore, this invention is directly related to artificial intelligence with inference abilities.

Title: CONTROLLING DISPLAY OF CONTENT

Patent Application Number: 20180253739
Patent Applicant: Google Inc.

Abstract: In one aspect, an automated endorsement prompt systems includes an endorsement prompt module comprising an endorsement signal module for retrieving an endorsement signal from an endorsement server; a search result module for retrieving search results from a search engine; a web history module for retrieving a web history for a user; and combiner logic for providing search results and an endorsement prompt. The combiner logic generates the endorsement prompt from the endorsement signal and the web history. A method for automatically generating endorsement prompts including the steps of: receiving a query from a user; obtaining additional information signals; obtaining a search result using the query; determining whether prompt behavior exists using the additional information; generating a prompt for an endorsement if the prompt behavior exists; and providing the search result and the prompt for presentation.

Explanation from LLM:
NO. The described system involves the use of various modules (endorsement signal module, search result module, web history module) and combiner logic to generate an endorsement prompt based on user activity and signals from an endorsement server. While this system uses logic and data processing to deliver personalized content, it does not explicitly involve advanced AI techniques such as machine learning, neural networks, or deep learning that would enable it to learn from its interactions or make inferences beyond the predefined rules. It operates more like a rule-based system that combines different inputs to produce an output, rather than an AI system with inference abilities.

Decision-making:

Query: *Is the following invention directly related to artificial intelligence that has decision-making ability? Respond with YES or NO and give an explanation.*

<u>Title: SYSTEM FOR FACILITATING BENEFACTION</u>
Patent Application Number: 20210174402
Patent Applicant: Wells Fargo Bank, N.A.

Abstract: A system for facilitating benefaction that identifies donation opportunities by monitoring and learning from users' behavior and suggesting donations based on the behavior of potential benefactors and potential beneficiaries. In certain aspects, the system cross-references donation conditions of one system user with donation conditions of one or more other system users to identify donation triggering events corresponding to donation opportunities.

Explanation from LLM:
YES. This invention is directly related to artificial intelligence that has decision-making ability because it involves a system that monitors and learns from users' behavior to suggest donations. The process of identifying donation opportunities through the analysis of user behavior and conditions, as well as suggesting donations based on this analysis, indicates the use of machine learning algorithms. These algorithms enable the system to make decisions about which donation opportunities to present to users, thereby incorporating decision-making capabilities typical of AI systems.

Title: NETWORK DETERMINATION OF LIMITED-ANGLE RECONSTRUCTION
Patent Application Number: 20230162412
Patent Applicant: Siemens Medical Solutions USA, Inc.

Abstract: A system and method include training of an artificial neural network to generate an output three-dimensional image volume based on input two-dimensional projection images, the training based on a plurality of subsets of two-dimensional projection images of each of a plurality of sets of two-dimensional projection images and associated ones of three-dimensional image volumes reconstructed from each of the plurality of sets of two-dimensional projection images.

Explanation from LLM:
NO. The described invention is related to machine learning, specifically in the context of training an artificial neural network to generate 3D images from 2D projections. While this involves sophisticated algorithms that can be part of artificial intelligence, it does not explicitly mention decision-making abilities as a core functionality. Instead, it focuses on image reconstruction and generation, which is more about pattern recognition and data transformation rather than making decisions based on those patterns or data.

Learning:

Query: *Is the following invention directly related to artificial intelligence that has learning ability? Respond with YES or NO and give an explanation.*

Title: ADAPTIVE INTERACTIVE MEDIA METHOD AND SYSTEM
Patent Application Number: 20190332656
Patent Applicant: Sunshine Partners, LLC

Abstract: An automated adaptive engine alters content in real-time based on feedback received. With this engine human adjustment is no longer necessary or is kept to minimum by choice. The engine may build an Emotional Profile (EP) from scratch, may modulate an existing EP based on detected real-time responses, or also may randomly try different things to detect new response from the user. The engine may be applied with an interactive media application, preferably embodied as an interactive book, that adjusts content automatically in real-time, based on a reader's quality of emotional response (or mental response) against expectation, without requiring any human intervention. The response feedback is detected as expressed through voice, facial expressions, vitals such as pulse or blood pressure, and/or activities in different parts of the brain, or other means of such expression.

Explanation from LLM:
YES. The described invention involves an automated adaptive engine that alters content in real-time based on feedback received, which implies a form of machine learning. This engine builds or modifies an Emotional Profile (EP) based on various forms of feedback, including voice, facial expressions, and physiological signals. It autonomously adapts its behavior without continuous human intervention, suggesting it uses some form of algorithmic learning to improve its performance over time. This capability aligns with the principles of artificial intelligence, particularly in the context of adaptive and responsive systems that can learn from interactions. Therefore, this invention is directly related to AI with learning abilities.

<u>Title: SYSTEMS AND METHODS FOR PROVIDING RECOMMENDATIONS BASED ON TRACKED ACTIVITIES</u>
Patent Application Number: 20220191568
Patent Applicant: OrCam Technologies Ltd.

Abstract:
A wearable apparatus is provided for capturing and processing images from an environment of a user. In one implementation, a wearable apparatus for monitoring activities includes a wearable image sensor configured to capture a plurality of images from an environment of a user of the wearable apparatus. The wearable apparatus also includes at least one processing device programmed to analyze the plurality of images to identify in one or more of the plurality of images at least one indicator of an activity, and to transmit, to an external device, the at least one indicator of the activity.

Explanation from LLM:
NO. While the described invention involves capturing and processing images which can be considered a form of data processing, it does not explicitly state that the apparatus has a learning ability or uses machine learning techniques to improve its performance over time. The processing device is programmed to analyze images and identify indicators of activities, but this alone does not imply the use of artificial intelligence with learning capabilities. For it to be directly related to AI with learning ability, the description would need to include aspects such as training the system on datasets, improving accuracy over time through experience, or using algorithms that learn from data to make decisions or predictions.

Creativity:

Query: *Is the following invention directly related to artificial intelligence that can generate content? Respond with YES or NO and give an explanation.*

Title: ADAPTIVE INTERACTIVE MEDIA METHOD AND SYSTEM
Patent Application Number: 20230004937
Patent Applicant: Sunshine Partners, LLC

Abstract: An AI response system configured to generate a response to an inquiry. The AI response system includes a learning unit and a response generation unit. The learning unit is machine-learned by an inquiry, general information, a response and a satisfaction level index for a response. The response generation unit is configured, based on a learning result of the learning unit, to generate, in response to inquiries from a plurality of mobile terminals, responses to be transmitted to each of the mobile terminals.

Explanation from LLM:
YES. The described invention involves an AI system designed to generate responses to inquiries, which aligns with the concept of AI-generated content. Specifically, the system includes a learning component that uses machine learning to understand and improve its responses based on past inquiries, general information, given responses, and satisfaction levels. This indicates the system is capable of evolving its content generation capabilities over time. Additionally, the response generation unit utilizes the learned data to produce tailored responses to new inquiries, clearly indicating its role in generating content dynamically.

Title: JOINT USE OF FACE, MOTION, AND UPPER-BODY DETECTION IN GROUP FRAMING
Patent Application Number: 20220006974
Patent Applicant: Polycom, Inc.

Abstract: A videoconferencing endpoint is described that uses a combination of face detection, motion detection, and upper body detection for selecting participants of a videoconference for group framing. Motion detection is used to remove fake faces as well as to detect motion in regions around detected faces during postprocessing. Upper body detection is used in conjunction with the motion detection in postprocessing to allow saving faces that have been initially detected by face detection for group framing even if the participant has turned away from the camera, allowing the endpoint to keep tracking the participants region better than would be possible based only on an unstable result coming from face detection.

Explanation from LLM:

NO. The invention described here is more closely related to computer vision techniques applied to improve the functionality of videoconferencing systems rather than directly related to artificial intelligence that generates content. The technologies mentioned (face detection, motion detection, upper body detection) are primarily used for improving the user experience by automatically adjusting the framing of participants in a video conference. While these techniques do involve AI (specifically machine learning models trained to recognize faces and body parts), they are not focused on generating new content like text, images, or videos. Instead, they enhance the existing visual content by optimizing how participants are displayed within the video conference frame.

Appendix B. Additional Descriptive Statistics of Skill-Technology Mapping

Table B1. AI-exposed occupations and required important skills from O*NET

This table reports the top 3 most-related occupations (3-digit SOC classification) for each type of AI technology and their corresponding required skills from the O*NET database. We identified the-most related occupations by sorting the number of required skills (weighted by the importance from the O*NET database) that mapped to each AI category in each occupation.

Technology category	Most-related occupation		2nd-most-related Occupation		3nd-most-related Occupation		Top 3 Related skills
	3-digit SOC	Title	3-digit SOC	Title	3-digit SOC	Title	Description
Language AI	491	Supervisors of Installation, Maintenance, and Repair Workers	492	Electrical and Electronic Equipment Mechanics, Installers, and Repairers	535	Water Transportation Workers	Quality Control Analysis, Reading Comprehension, Monitoring
Perception AI	491	Supervisors of Installation, Maintenance, and Repair Workers	492	Electrical and Electronic Equipment Mechanics, Installers, and Repairers	493	Vehicle and Mobile Equipment Mechanics, Installers, and Repairers	Repairing; Monitoring; Active Listening
Engagement AI	331	Supervisors of Protective Service Workers	212	Religious Workers	431	Supervisors of Office and Administrative Support Workers	Coordination; Social Perceptiveness; Active Listening

Inference AI	172	Engineers	451	Supervisors of Farming, Fishing, and Forestry Workers	491	Supervisors of Installation, Maintenance, and Repair Workers	Critical Thinking; Reading Comprehension; Monitoring
Decision-making AI	451	Supervisors of Farming, Fishing, and Forestry Workers	491	Supervisors of Installation, Maintenance, and Repair Workers	535	Water Transportation Workers	Coordination; Operation and Control; Management of Personnel Resources
Learning AI	111	Top Executives	331	Supervisors of Protective Service Workers	532	Air Transportation Workers	Active Listening; Complex Problem Solving; Critical Thinking
Creativity AI	111	Top Executives	113	Operations Specialties Managers	431	Supervisors of Office and Administrative Support Workers	Complex Problem Solving; Critical Thinking; Speaking

Table B2. Examples of individual skill sets from Revelio Labs

This table shows individual-level skill sets for examples of newly-hired workers and departed workers. The individual-level skills are obtained from the Revelio Labs database (collected from individual online profiles).

Panel A. Individual skill sets from new hires

Job Hirings					
Firm	Year	Occupation	Within firm transition	User_id	Skill set
STARBUCKS CORP	2014	Supervisors of Sales Workers (411)	No	49262072	call centers, customer satisfaction, customer service, insurance, inventory management, leadership, management, merchandising, Microsoft excel, Microsoft office, Microsoft word, PowerPoint, process improvement, retail, sales, team building, team leadership, time management, training, visual merchandising
NATIONAL INSTRUMENTS CORP	2015	Engineers (172)	Yes	49286590	analog, application-specific integrated circuits (ASIC), automation, C, C++, data acquisition, debugging, electrical engineering, electronics, embedded software, embedded systems, engineering, field-programmable gate arrays (FPGA), integrated circuits (IC), LabVIEW, manufacturing, MATLAB, PCB design, product development, product management, project management, semiconductors, soc, software tests, system design, test automation, test engineering, test equipment, testing
ROYALITE PETROLEUM CO INC	2014	Media and Communication Workers (273)	No	49267198	advertising, business development, business management, business strategy, creativity, customer relationship management (CRM), digital marketing, e-commerce, leadership, management, marketing, marketing strategy, Microsoft office, mobile devices, negotiation, online advertising, online marketing, product management, project management, public speaking, research, social media, social networks, Spanish, strategic planning, strategy, team leadership, teamwork, telecommunications, training

Panel B. Individual skill sets from job separations

Job Separations					
Firm	Year	Occupation	Within firm transition	User_id	Skillset
VMWARE INC-CL A	2010	Computer Occupations (151)	Yes	49282976	cisco technologies, cloud computing, data center, disaster recovery, information technology infrastructure library, integration, IT service management, Linux, network security, networking, requirements analysis, servers, software development, storage area networks, structured query language, TCP/IP, UNIX, virtualization, VMware, windows server
STARBUCKS CORP	2016	Supervisors of Food Preparation and Serving Workers (351)	No	49250337	catering, cooking, culinary skills, customer service, event planning, food, food & beverage, food safety, food service, hospitality, hospitality management, inventory management, leadership, management, menu development, Microsoft excel, Microsoft office, Microsoft word, PowerPoint, process improvement, project management, research, restaurant management, restaurants, sales, team building, time management, training
WALMART INC	2020	Other Transportation Workers (536)	No	49285556	cashiering, customer satisfaction, customer service, data entry, inventory control, inventory management, leadership, loss prevention, management, merchandising, Microsoft excel, Microsoft office, Microsoft word, PowerPoint, retail, retail sales, sales, team building, teamwork, time management

Appendix C. Additional Analysis

Table C1. AI innovation and firm employment - 2SLS First-stage regressions

This table reports the first-stage results of 2SLS regressions examining the effects of AI innovation on firm outcomes. The dependent variables are *Log # of "type" patents* interacted with the *Exposed* indicator, which equals one if the occupation requires any important (importance score >= 4) skills that are related to the AI in the given type (See Section 4.1). The instruments include (1) the total patent examiner leniency for the firm's AI patents in the given type during the year interacted with the *Exposed* dummy (2) the total patent examiner leniency for the firm's other patents in during the year interacted with the *Exposed* dummy and (3) the total patent examiner leniency for the firm's all patent applications filed during the year interacted with the *Not Exposed* dummy. Control variables are defined in Table 4. Each regression includes firm-level control variables in year t-1, patent applications-by-year fixed effects corresponding to the numbers of different "type" and "other" patent applications filed by the firm, and firm fixed effects. All continuous variables are winsorized at 1% and 99%. T-statistics are reported in parentheses. *, **, and *** denote significance at 10%, 5%, and 1%.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Dependent Variable						
	Exposed × Log # of Language Patents	*Exposed × Log # of Perception Patents*	*Exposed × Log # of Engagement Patents*	*Exposed × Log # of Inference Patents*	*Exposed × Log # of Decision-Making Patents*	*Exposed × Log # of Learning Patents*	*Exposed × Log # of Creativity Patents*
IV: Examiner leniency for type *k* AI patents × Exposed	0.817***	0.260***	0.384***	0.1979***	0.218***	0.288***	0.770***
	(23.13)	(9.20)	(17.98)	(12.51)	(11.59)	(12.06)	(13.90)
IV: Examiner leniency for non-type-*k* patents × Exposed	0.004***	0.010***	0.003***	0.0039**	0.009***	0.0029**	0.007***
	(7.31)	(9.74)	(4.05)	(2.51)	(5.32)	(2.54)	(12.47)
IV: Examiner leniency for all patents × Non-Exposed	-0.011***	-0.020***	-0.020***	-0.027***	-0.022***	-0.026***	-0.008***
	(-19.07)	(-21.18)	(-27.87)	(-21.73)	(-17.52)	(-28.43)	(-16.34)
Exposed	0.145***	0.256***	0.312***	0.560***	0.511***	0.408***	0.092***
	(27.19)	(42.37)	(37.4)	(55.67)	(56.28)	(46.55)	(20.62)
Log Assets	0.002	0.007***	0.006**	0.014***	0.011***	0.006**	-0.002*
	(1.51)	(2.74)	(2.47)	(4.50)	(4.53)	(2.41)	(-1.75)
ROA	-0.0014	0.014*	0.015**	0.008	0.011	0.009	0.011***
	(-0.26)	(1.73)	(2.18)	(0.86)	(1.53)	(1.10)	(3.07)

MTB	0.0003	0.0006	0.001	0.0009	0.001	0.0004	0.0001
	(0.61)	(0.84)	(1.59)	(1.08)	(1.44)	(0.51)	(0.13)
R&D	-0.002	0.0003	0.003	-0.0045	-0.0008	0.0021	0.0005
	(-1.38)	(0.13)	(1.38)	(-1.60)	(-0.35)	(0.84)	(0.52)
R&D missing	-0.012	-0.0071	0.0014	-0.023	-0.013	0.0014	-0.0003
	(-1.63)	(-0.78)	(0.11)	(-1.62)	(-1.16)	(0.10)	(-0.05)
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
# of AI type *k* apps × year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
# of other apps × year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
# all apps × year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Sanderson-Windmeijer *F*	1,251.7	1,127.7	1,145.4	1,087.2	1,048.6	1,117.6	832.4
Observations	703,031	703.031	703,031	703,031	703,031	703,031	703,031

Table C2. AI innovation and firm value - 2SLS first-stage regression

This table reports the first-stage results of 2SLS regressions examining the effects of AI innovation on firm outcomes. The dependent variables include Log # of augmenting AI patents, Log # of scope augmenting AI patents, Log # of core augmenting AI patents, and Log # of displacing AI patents. *Log # of augmenting (displacing) AI patents* is the log-transformed number of patent grants in AI categories that have a positive (negative) association with employment growth. *Log # of scope (core) augmenting* AI patents is the log-transformed number of patent grants in AI categories that have a positive association with overall employment growth in headcount and have (do not have) a positive association with employment growth among workers with skills that are unique relative to their occupation median. The instruments include (1) total patent examiner leniency for the firm's AI patents in the given type during the year and (2) the total patent examiner leniency for the firm's other patent applications filed during the year. Control variables are defined as in Table 4. Each regression includes firm-level control variables in year t-1, patent applications-by-year fixed effects corresponding to the numbers of different "types" and "other" patent applications filed by the firm, and firm fixed effects. All continuous variables are winsorized at 1% and 99%. T-statistics are reported in parentheses. *, **, and *** denote significance at 10%, 5%, and 1%.

	(1)	(2)	(3)	(4)
VARIABLES	*Log # Augmenting*	*Log # Scope augmenting*	*Log # Core Augmenting*	*Log # Displacing*
IV: Examiner leniency for augmenting	0.115***			
	(14.065)			
IV: Examiner leniency for non-augmenting	0.001			
	(0.829)			
IV: Examiner leniency for scope augmenting		0.146***		
		(13.345)		
IV: Examiner leniency for non-scope augmenting		0.000		
		(0.584)		
IV: Examiner leniency for core augmenting			0.114***	
			(14.929)	
IV: Examiner leniency for non-core augmenting			0.001**	
			(2.259)	
IV: Examiner leniency for displacing				0.164***
				(11.088)
IV: Examiner leniency for non-displacing				0.000
				(0.809)

Size	0.007***	0.007***	0.007***	0.004
	(2.635)	(3.096)	(3.013)	(1.400)
MTB	0.000***	0.000***	0.000***	-0.000
	(3.562)	(3.389)	(5.172)	(0.271)
ROA	-0.003	0.002	-0.004	-0.000
	(0.544)	(0.283)	(0.645)	(0.022)
R&D	-0.003	-0.001	-0.004	-0.001
	(0.944)	(0.589)	(1.488)	(0.415)
R&D missing	-0.021	-0.017	-0.026*	-0.015**
	(1.360)	(1.634)	(1.923)	(2.275)
Firm fixed effects	Yes	Yes	Yes	Yes
# of AI type *k* apps × year FEs	Yes	Yes	Yes	Yes
# of other apps × year FEs	Yes	Yes	Yes	Yes
Sanderson-Windmeijer *F*	100.03	129.10	110.37	122.16
Observations	14,299	14,309	14,310	14,334

Figure 1. Growth in AI innovation activity over time

This figure illustrates the time series of AI patenting activity by U.S. public firms from June 2007 to December 2023. Patent grants are counted as of the publication date of grant. Patent data are from the USPTO Bulk Data Storage System (BDSS). Panel A shows the time series for AI patents and for all patents, while Panel B shows the time series for each AI patent category. Patent frequencies are cumulative and shown on a monthly basis, scaled by the patent count in June 2007 (as the base group). AI patent categories are defined according to the typology outlined in Section 2, and individual patents are categorized according to inferences from a Large Language Model (LLM) as described in Section 3.2.

Panel A. AI innovation and overall innovation

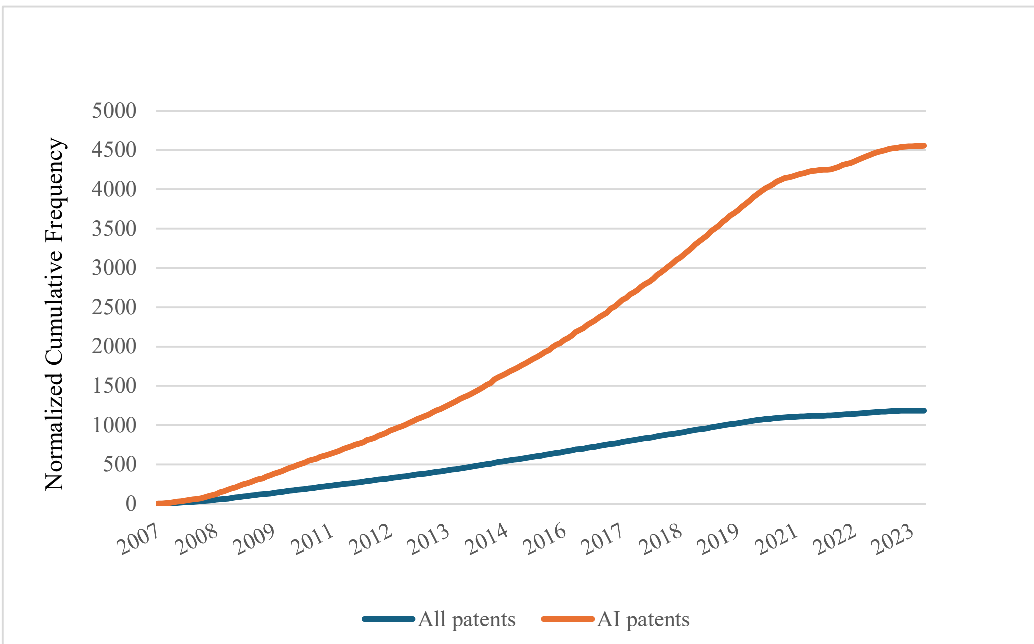

Panel B. AI innovation, by category

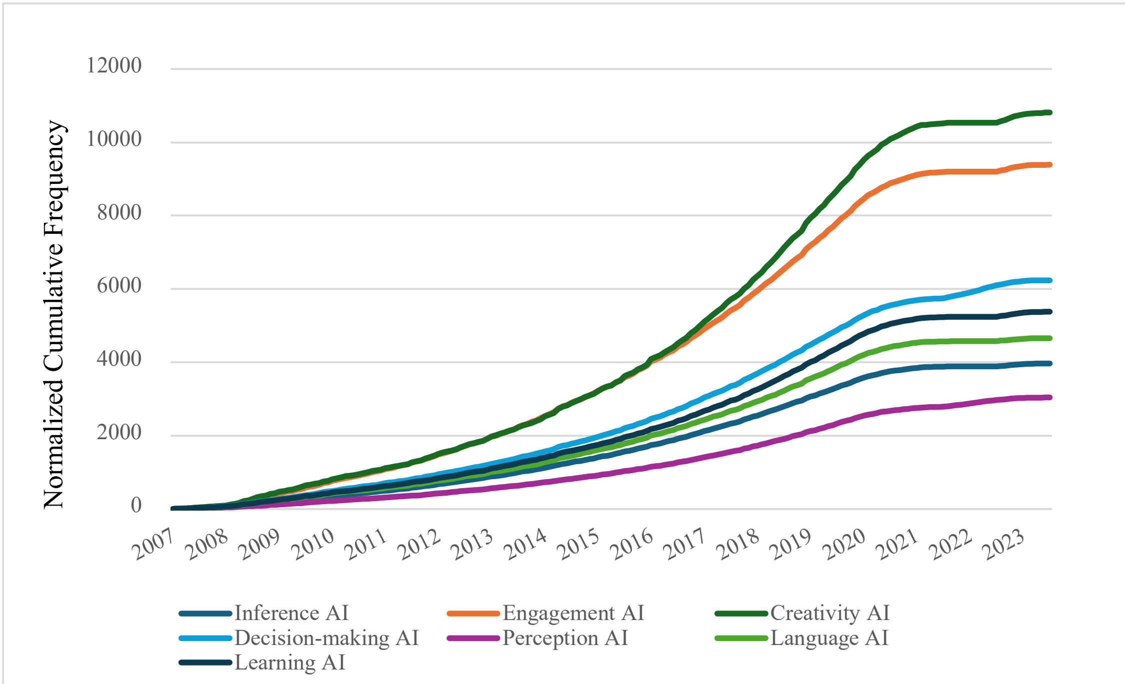

Figure 2. AI innovation activity and firm characteristics

This figure illustrates the extent of different types of AI innovation by different categories of firms over 2007-2023. Statistics are summarized as of the year of the publication date of patent grants. Patent data are from the USPTO Bulk Data Storage System (BDSS). Panel A reports, by technology category, AI patenting as a fraction of total patents granted to small and large firms in the CRSP/Compustat database. A firm is large (small) if the firm's log-transformed total assets is above (below) the sample median in a given year. Panel B reports, by technology category, AI patenting as a fraction of total patents granted to labor-intensive and capital-intensive firms. A firm is defined as labor-intensive (capital-intensive) if the firm's total employment scaled by total assets is above (below) the sample median in a given year. Panel C reports, by technology category, AI patenting as a fraction of total patents granted to R&D-intensive and Non-R&D-intensive firms. A firm is R&D-intensive (non-R&D-intensive) if the firm's R&D expenditures scaled by total assets is above (below) the sample median in a given year. Firm employment data are obtained from the Revelio Labs database. Other firm-level data are from Compustat.

Panel A. AI innovation activity in small vs. large firms

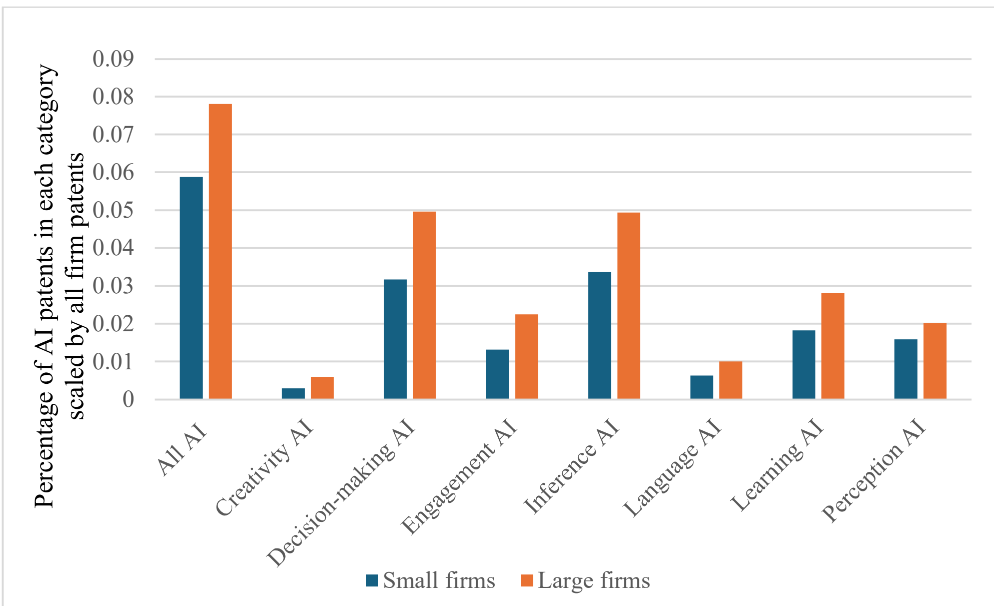

Panel B. AI Innovation in labor-intensive vs. capital-intensive firms

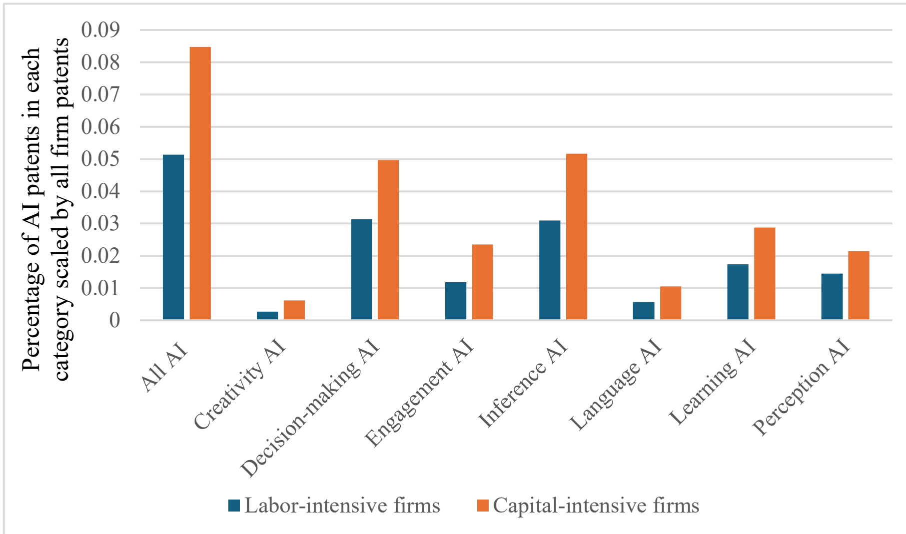

Panel C. AI innovation in R&D intensive vs. non-R&D intensive firms

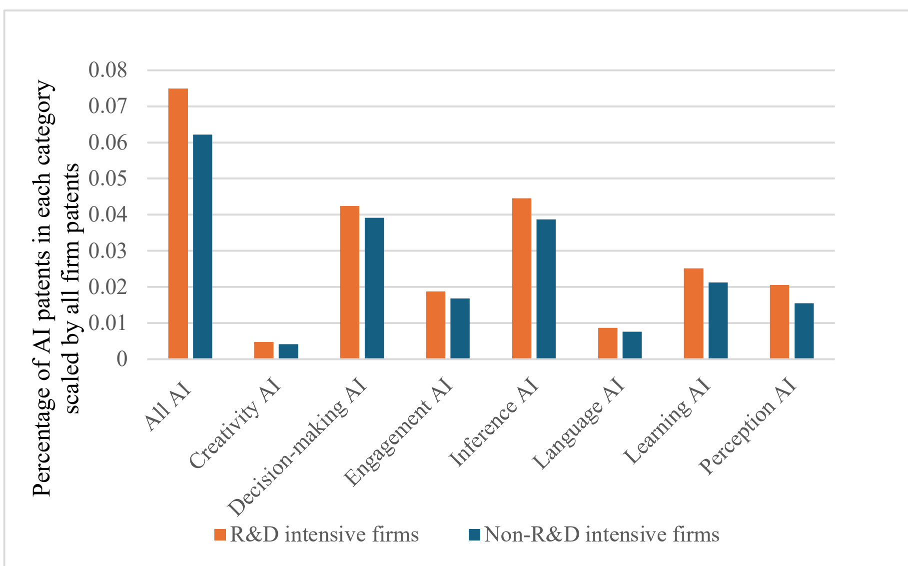

Figure 3. Occupational employment effects of AI innovation

This figure shows heatmaps illustrating the effects of AI innovation on occupational employment growth. For each AI category, we estimate the effects of AI innovation by running 85 regressions (one for each 3-digit SOC occupation in a given firm-year). We assign the occupations with significant positive estimates as 1, negative estimates as -1, and insignificant estimates as 0. Assigned scores at the 3-digit SOC level are then summed to the 2-digit SOC level, with weights corresponding the employment share of each 3-digit SOC occupation. Panel A shows a heatmap for all occupations, and Panel B shows a heatmap for occupations with exposure to AI as defined in Section 4.1. The employment data for each O*NET occupation are obtained from the Revelio Labs database.

Panel A. All occupations

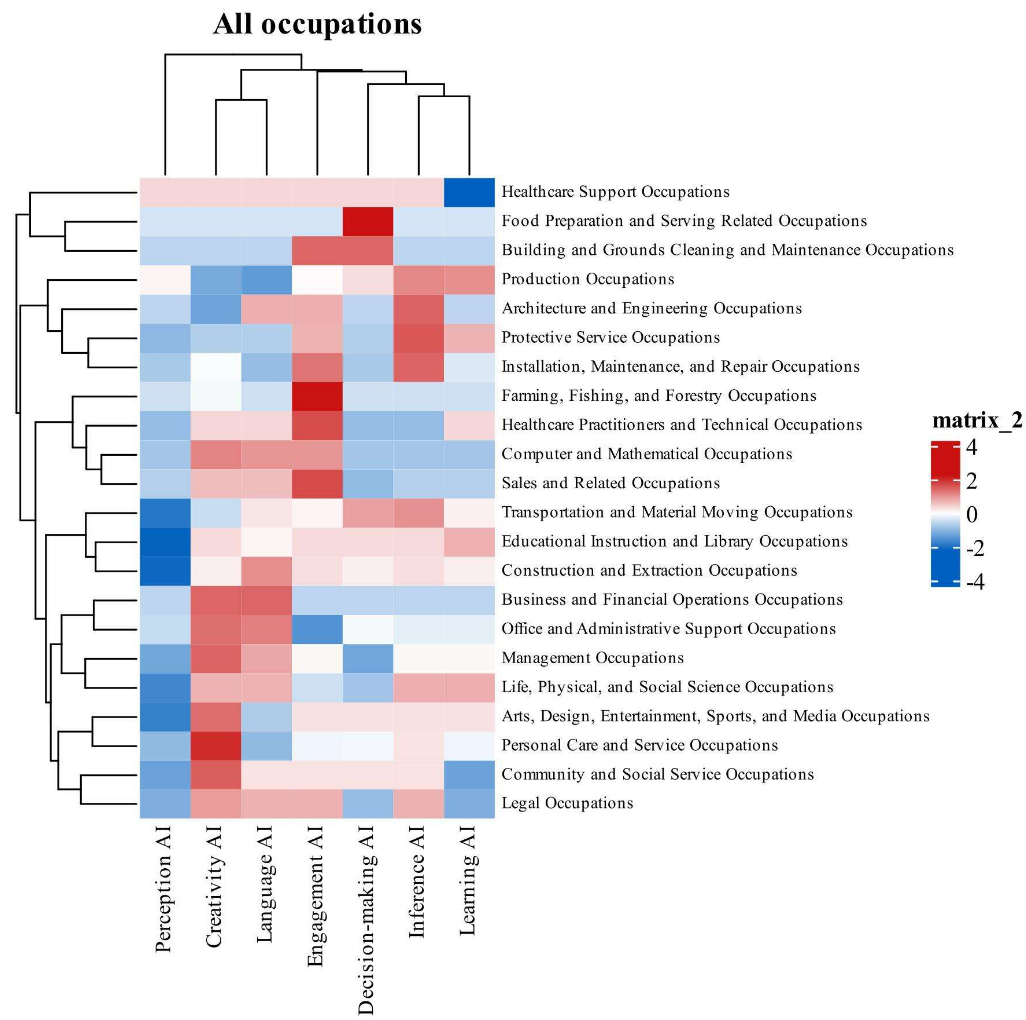

Panel B. Exposed occupations

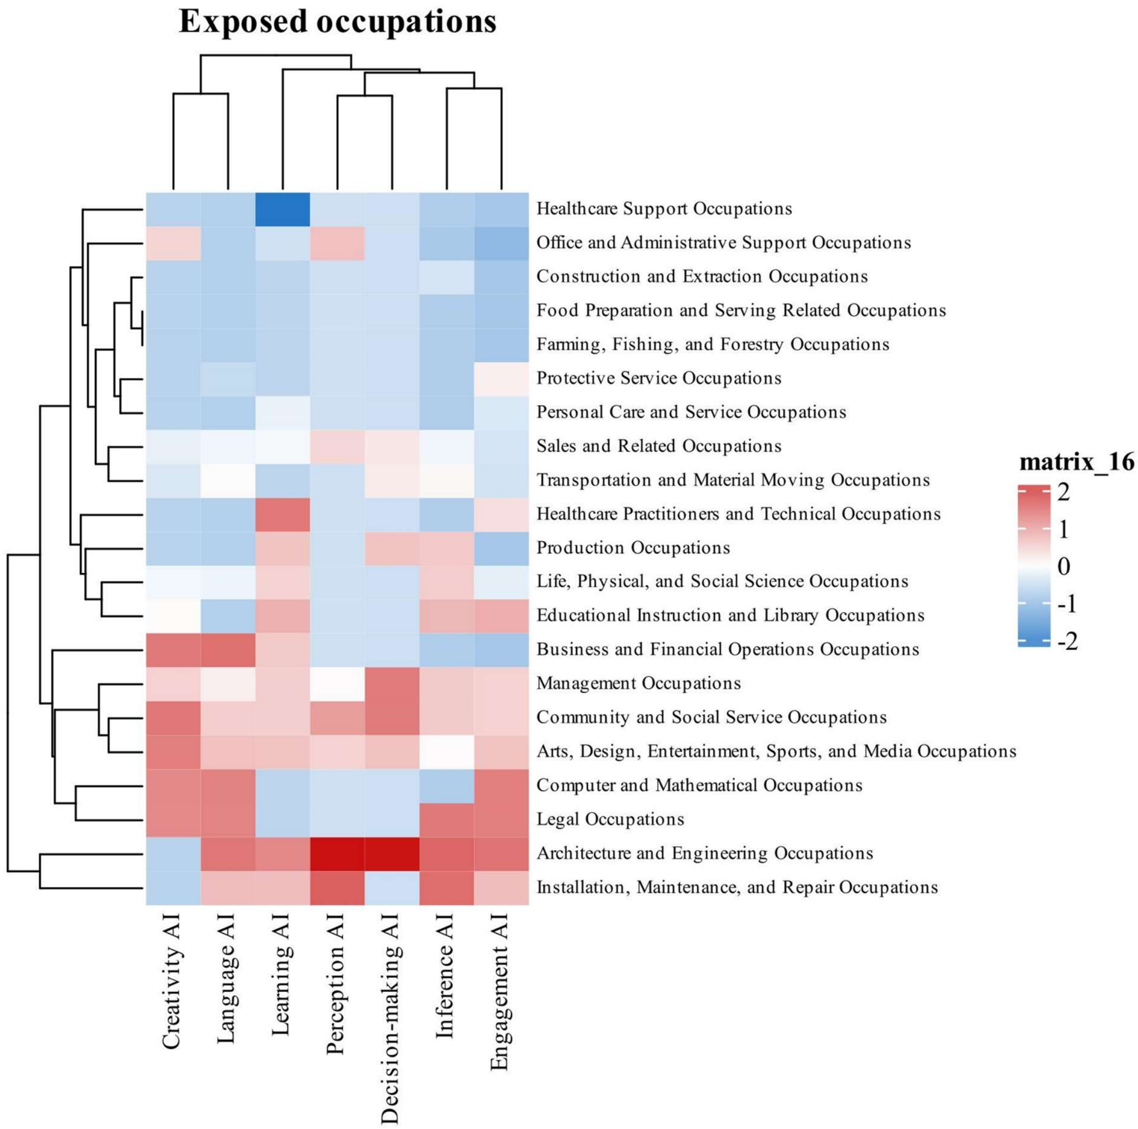

Table 1. Domains of human cognition and functional capabilities of artificial intelligence

This table provides a typology of AI based on functional capabilities. The first column lists the six principal domains of cognition specified by Newell (1990). The second column lists the six key domains of neurocognition described in the Diagnostic and Statistical Manual of Mental Disorders (DSM-5), along with relevant subdomains for each. The third column shows seven capabilities, each of which is a leading example of AI functionality within a cognitive domain.

Principal Area of Human Cognition (Newell, 1990)	Key Domain of Neurocognitive Function (DSM-5, 2013)	AI Capability
Language	**Language:** expressive language (including naming, word finding, fluency, grammar, and syntax) and receptive language	*Language*
Perception, motor behavior	**Perceptual-motor:** includes abilities subsumed under the terms visual perception, visuoconstructional, perceptual-motor, praxis, and gnosis	*Perception/Motor*
Motivation, emotion	**Social cognition:** recognition of emotions, theory of mind	*Engagement*
Problem solving, decision making, routine action	**Complex attention:** sustained attention, divided attention, selective attention, processing speed	*Inference*
	Executive function: planning, decision-making, working memory, responding to feedback/error correction, overriding habits/inhibition, mental flexibility	*Decision making*
Memory, learning, skill	**Learning and memory:** immediate memory, recent memory (including free recall, cued recall, and recognition memory), very long-term memory (semantic; autobiographical), implicit learning	*Learning*
Imagining, dreaming, daydreaming		*Creativity*

Table 2. Tasks and examples corresponding to different AI types

This table shows examples of specific tasks and real-world applications related to the seven AI types defined in Section 2. Examples of AI Applications are collected from various public sources including news articles, research papers, industry reports, etc.

AI Type	Tasks	Examples of AI Applications
Language	Text generation; natural language processing; voice synthesis; language translation and interpretation	Large language models (LLMs); dubbing for video localization; voiceprint analysis; real-time voice translation; live call-center interpretation
Perception/Motor	Image, sound, and speech recognition; object detection; sensor fusion; tactile (touch) perception	Computer vision systems; Autonomous vehicle systems; Diagnostic medical imaging; Augmented reality (AR) and Virtual reality (VR)
Engagement	Interactive chat; interactive content generation; customization; emotion detection; context-aware services	Siri; Cortana; Alexa; Chat GPT; customer service AI; videogame NPCs; smart home devices; self-driving delivery robots
Inference	Logical reasoning; heuristic inference; optimization; prediction	AlphaFold protein structure prediction; fraud detection systems; expert systems; automated reasoning systems; C++ coding assistants; robo-advisors
Decision-making	Rule-based decisioning; learning-based decision-making; goal-based planning	Automated credit decisioning; AI resume screening; autonomous vehicle navigation
Learning	Machine learning; reinforcement learning; deep learning	Personalized recommendation engines; machine learning systems for healthcare diagnostics, anomaly detection, and autonomous vehicles
Creativity	Text, image, sound, video, and art generation; procedural content generation; transformation	ChatGPT; Adobe Firefly art generator; OpenAI Sora video generator; robotic painters and sculptors

Table 3. Artificial intelligence innovation activity, by industry and IPC class

This table summarizes the frequencies of AI patents granted across industrial sectors (Panel A) and across major IPC classes (Panel B) for each category of AI innovation during 2007-2023. Statistics are summarized as of the year of the publication date of patent grants. Patent data are from the USPTO Bulk Data Storage System (BDSS). The sample consists of patent grants in the USPTO bulk data that are filed by companies in CRSP/Compustat.

Panel A. Industry distribution of AI patent grants

Description	SIC (one-digit)	All AI	Language	Perception	Engagement	Inference	Decision-making	Learning	Creativity
Manufacturing	3	67,285	22,530	22,886	18,336	38,836	4,771	8,552	42,189
Transport; storage and communication	7	48,008	12,968	20,409	19,892	28,373	6,836	11,591	32,266
Electricity; gas and water supply	4	9,706	2,922	3,580	3,377	5,505	1,047	1,733	6,034
Mining and quarrying	2	4,850	1,629	1,729	720	2,743	190	238	2,986
Wholesale and retail trade	6	4,705	1,165	1,672	1,401	3,016	335	620	2,733
Community; social and personal services	9	4,545	1,572	1,639	661	2,669	149	277	3,081
Construction	5	2,080	592	826	914	1,206	233	417	1,429
Financial; insurance; real estate and business services	8	1,475	470	729	622	1,016	189	341	1,138
Agriculture; hunting; forestry and fishing	1	1,441	406	565	168	824	48	63	936
Others	0	30	11	6	2	13	0	0	12

Panel B. Key IPC class distribution

Description	IPC (3-digit)	All AI	Language	Perception	Engagement	Inference	Decision-making	Learning	Creativity
Computing, Calculating, or Counting	G06	76,053	13,771	24,496	27,785	50,452	43,874	30,631	9,104
Electric Communication Technique	H04	32,123	4,024	8,685	9,701	18,012	18,024	9,719	2,055
Musical Instruments; Acoustics	G10	7,581	4,967	3,817	5,166	5,786	3,963	3,699	1,623
Measuring; Testing	G01	6,394	210	2,248	901	4,330	3,906	1,937	77
Medical or Veterinary Science; Hygiene	A61	6,330	144	2,317	836	3,967	3,602	1,926	61
Controlling; Regulating	G05	5,365	195	2,261	1,148	3,943	4,538	2,050	68
Vehicles in General	B60	4,967	156	2,647	993	3,582	4,274	1,536	30
Signaling	G08	3,962	163	1,849	856	2,898	3,241	1,167	31
Educating; Cryptography; Display; Advertising; Sales	G09	2,104	359	953	1,026	1,256	1,014	694	234
Sports; Games; Amusements	A63	1,266	132	490	692	755	674	380	115

Table 4. Characteristics of firms, AI Innovation, and employment

This table reports summary statistics for occupational employment growth and firm characteristics during the sample period 2007-2023. Firm employment data are obtained from the Revelio Labs database. In Panel A, *Employment growth (with/without new skills)* is measured as the changes in employees (with/without new skills) in an occupation scaled by the total number of employees within the firm. For ease of exposition, the growth rate is multiplied by a factor of 1,000. *# Hirings (Job separations) with (without) new skills* is the number of employees hired (departed) with (without) new skills. We define employees with new skills as those with skills that are absent from the occupational median skill set. The individual skills data are obtained from Revelio Labs. Panel B reports the firm AI innovation activities and other characteristics. Patents in each AI category are identified from the LLM and defined in Section 2. Patents are defined as *Augmenting (displacing) AI* if, according to Table 5, Panel A, the given category of AI innovation has a positive (negative) estimated effect on subsequent labor growth in exposed occupations. Since the labor displacement effects are partially driven by the number of displaceable workers that a firm has, we further require displacing AI patents to be related to occupations that account for more than half of the total employees of the firm. Patents are defined as *Scope augmenting AI* if, according to Table 5, Panel B, the given category of AI innovation has a positive estimated effect on the growth of employees with new skills. Patents are defined as *Core augmenting AI* if, according to Table 5, Panels A and B, the given category of AI innovation has a positive estimated effects only on the firm's overall labor growth but not on the growth of employees with new skills. *Size* is the natural log of total assets in the prior year from Compustat. *MTB* is the market-to-book ratio in the prior year. *ROA* is income before extraordinary items divided by total assets in the prior year. *R&D* is the natural log of R&D expenditures in the prior year. Missing values of R&D are imputed as zero and indicated by *R&D Missing*. *Total Factor Productivity* $_{t,\,t+1}$ is the log-transformed change in yearly total factor productivity (TFP). The firm-level TFP is the estimated residual from a firm-year level linear production function model with capital and labor inputs. Firm-level production inputs in the residual estimation regression are proxied for by the cost of goods sold and the firm's total employment. The dependent variable in the residual estimation regression is the firm production output measured as net sales. Residual estimates are scaled in thousands before the log transformation. *Operating expenses* $_{t,\,t+1}$ is measured as the yearly growth rate of total operating expenses (Compustat item XOPR). *SG&A* $_{t,\,t+1}$ is measured as the yearly growth rate of selling, general, and administrative expenditures (Compustat item XSGA). *Tobin`s Q* $_{t,\,t+1}$ is measured as the yearly changes in the ratio of the book value of debt (Compustat items *DLTT* + *DLC*) plus the market value of equity (Compustat items *PRCC_F* × *CSHO*) minus the firm's current assets (Compustat item *ACT*) to the book value of property, plant, and equipment (Compustat item *PPEGT*). All other variables are as described in Section 3.

Panel A. Occupation-level employment characteristics

	N	Mean	SD	p10	Median	p90
Employment growth	133,482	-2.352	17.674	-15.535	-0.107	8.827
Employment growth with new skills	136,969	0.006	0.684	-0.489	0.000	0.520
Employment growth without new skills	136,969	-0.090	0.389	-0.108	0.000	0.000
# Hirings with new skills	136,969	4673.386	10253.119	0.000	1000	11000
# Hirings without new skills	136,969	178.639	410.434	0.000	0.0	1000
# Job separations with new skills	136,969	23.633	151.904	0.000	0.0	0.0
# Job separations without new skills	136,969	4472.377	9177.034	0.000	1000	11000

Panel B. Firm-level characteristics

AI innovation	N	Mean	SD	p10	Median	p90
# All AI	13,710	7.703	29.138	0.000	0.000	13.000
# Language AI	13,710	0.729	3.437	0.000	0.000	1.000
# Perception AI	13,710	2.235	9.502	0.000	0.000	3.000
# Engagement AI	13,710	1.773	7.829	0.000	0.000	2.000
# Inference AI	13,710	4.650	18.459	0.000	0.000	8.000
# Decision-making AI	13,710	4.020	16.175	0.000	0.000	7.000
# Learning AI	13,710	2.589	10.283	0.000	0.000	4.000
# Creativity AI	13,710	0.447	2.186	0.000	0.000	0.000
# Augmenting AI	13,710	4.835	19.498	0.000	0.000	8.000
# Scope-augmenting AI	13,710	3.321	13.616	0.000	0.000	5.000
# Core-augmenting AI	13,710	4.240	17.141	0.000	0.000	7.000
# Displacing AI	13,710	2.214	9.788	0.000	0.000	2.000
Firm Characteristics						
Size	13,710	6.810	2.437	3.659	6.706	10.177
MTB	13,213	2.525	2.006	1.009	1.839	4.865
ROA	13,697	-0.029	0.354	-0.443	0.090	0.204
R&D	13,710	3.282	2.133	0.000	3.382	6.016
R&D missing	13,710	0.142	0.349	0.000	0.000	1.000
Total Factor Productivity $_{t, t+1}$	11,818	0.045	0.360	-0.082	0.006	0.238
Operating expenses $_{t, t+1}$	12,712	0.106	0.340	-0.138	0.061	0.367
SG&A $_{t, t+1}$	10,982	0.088	0.275	-0.114	0.054	0.300
Tobin`s Q $_{t, t+1}$	12,193	-1.010	20.741	-7.569	0.048	5.852

Table 5. Effects of AI Innovation on employment growth

This table reports the second-stage results of 2SLS regressions examining the effects of AI innovation on employment growth at the firm-occupation-year level. First-stage results are reported in Appendix C. The endogenous variables in each regression are as follows: (1) *Log # of type k AI patents × Exposed*, the log-transformed number of granted patent applications in an AI category by the firm interacted with an indicator equal to one if the occupation requires any skills that are related to the given type of AI technology; (2) *Log # of other patents × Exposed*, the log-transformed number of all non-type granted patent applications by the firm, interacted with the *Exposed* indicator; (3) *Log # of all patents × Non-Exposed*, the log-transformed number of all granted patent applications by the firm, interacted with the *Non-Exposed* indicator. The instruments are reported and defined in Appendix C. In Panel A, the dependent variable is the *Employment growth (headcount)*, calculated as the change in employees in an occupation scaled by the total number of employees within the firm. In Panel B, the dependent variable is *Employment growth (with new skills)*, defined as the change in employees with new skills in an occupation scaled by the total number of employees within the firm. We define an employee as having new skills if his/her skill set contains skills that are absent from the occupational median skill set. The individual skill data are obtained from Revelio Labs. Control variables are as defined in Table 4. Each regression includes patent applications-by-year fixed effects corresponding to the numbers of different "type" and "other" patent applications filed by the firm, industry-by-year fixed effects (SIC 2-digit), and firm fixed effects. All continuous independent variables are winsorized at 1% and 99%. T-statistics are reported in parentheses. *, **, and *** denote significance at 10%, 5%, and 1%.

Panel A. Effects of AI innovations on headcount growth

	Employment growth (headcount)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	AI Type k =						
	Language	*Perception*	*Engagement*	*Inference*	*Decision-making*	*Learning*	*Creativity*
Log # of type k AI patents × Exposed	0.324***	-0.256**	0.354***	0.110	0.193**	0.390***	0.612***
	(4.29)	(-2.16)	(3.64)	(1.08)	(2.09)	(3.23)	(4.60)
Log # of other patents × Exposed	1.35***	-0.045	1.295***	-0.153	0.471	0.334	0.579
	(2.99)	(-0.11)	(2.96)	(-0.39)	(1.15)	(0.79)	(1.39)
Log # of all patents × Non-Exposed	1.355***	-0.116	1.561***	0.018	0.694	0.674	0.893**
	(2.99)	(-0.29)	(3.42)	(0.04)	(1.59)	(1.52)	(2.10)
Exposed	0.118	-0.721***	0.085	-0.106	0.162	0.271	0.250**
	(1.46)	(-4.05)	(0.63)	(-0.72)	(1.28)	(1.59)	(2.06)
Size	-0.321***	-0.296***	-0.326***	-0.314***	-0.317***	-0.304***	-0.316***
	(-4.38)	(-4.05)	(-4.43)	(-4.27)	(-4.34)	(-4.16)	(-4.32)
ROA	1.372***	1.396***	1.453***	1.399***	1.489***	1.472***	1.420***
	(3.57)	(3.62)	(3.77)	(3.62)	(3.86)	(3.79)	(3.70)

MTB	0.358***	0.360***	0.357***	0.371***	0.369***	0.359***	0.360***
	(12.54)	(12.56)	(12.43)	(12.98)	(12.91)	(12.65)	(12.64)
R&D	-0.944***	-0.966***	-0.953***	-0.951***	-0.954***	-0.962***	-0.967***
	(-13.68)	(-13.75)	(-13.64)	(-13.46)	(-13.74)	(-13.70)	(-13.84)
R&D missing	-3.28***	-3.33***	-3.36***	-3.288***	-3.308***	-3.341***	-3.318***
	(-10.88)	(-10.97)	(-11.13)	(-10.80)	(-11.02)	(-10.98)	(-10.86)
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
# of AI type *k* apps × year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
# of other apps × year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
# all apps × year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SIC-2 × year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Kleibergen-Paap F	119.5	590.9	578.3	578.7	549.3	573.1	563.3
Observations	703,031	703.031	703,031	703,031	703,031	703,031	703,031
R-squared	0.004	0.006	0.005	0.005	0.005	0.005	0.005

Panel B. Effects of AI innovation on new skill growth

	Employment growth (new skills)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	AI Patent Type *k* =						
	Language	*Perception*	*Engagement*	*Inference*	*Decision-Making*	*Learning*	*Creativity*
Log # of type *k* AI patents × Exposed	0.0004	-0.0019**	0.0019**	0.0001	-0.0007	0.0024***	0.0034***
	(0.66)	(-1.96)	(2.53)	(0.11)	(-0.96)	(2.60)	(3.01)
Log # of other patents × Exposed	-0.004*	-0.0023	-0.002	-0.0032	-0.003*	-0.0034*	-0.0042*
	(-1.86)	(-1.19)	(-0.90)	(-1.64)	(-1.66)	(-1.68)	(-1.88)
Log # of all patents × Non-Exposed	-0.0035	-0.003	-0.0003	-0.0027	-0.0036*	-0.0013	-0.0025
	(-1.64)	(-1.55)	(-0.10)	(-1.27)	(-1.66)	(-0.64)	(-1.10)
Exposed	0.0022***	0.0001	0.006***	0.003***	0.0012	0.007***	0.0063***
	(3.09)	(0.09)	(5.43)	(2.71)	(1.19)	(4.91)	(6.00)
Size	-0.0002	-0.0001	-0.0019	-0.0002	-0.0001	-0.0001	-0.0001
	(-0.69)	(-0.43)	(-0.72)	(-0.84)	(-0.60)	(-0.49)	(-0.31)
ROA	0.0035***	0.0031***	0.003***	0.004***	0.003***	0.0036***	0.0034***
	(5.30)	(4.72)	(5.04)	(5.56)	(4.58)	(5.33)	(5.05)
MTB	0.0001	0.00001	0.0001	0.0001	0.0001	0.00001	0.0001
	(0.98)	(1.24)	(0.67)	(1.37)	(1.16)	(0.84)	(0.81)
R&D	0.0005**	0.0005**	0.0001**	0.0005**	0.0001*	0.0005**	0.0006***
	(2.17)	(2.04)	(2.31)	(2.39)	(1.95)	(2.10)	(2.69)
R&D missing	-0.0009	-0.002*	-0.0011	-0.001	-0.0014	-0.0012	-0.0005
	(-0.87)	(-1.65)	(-1.04)	(-1.04)	(-1.30)	(-1.17)	(-0.48)
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
# of AI type *k* apps × year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
# of other apps × year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
# all apps × year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SIC-2 × year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Kleibergen-Paap F	537.5	590.9	578.3	578.7	549.3	573.1	563.3
Observations	703,031	703,031	703,031	703,031	703,031	703,031	703,031
R-squared	-0.0001	0.0004	-0.0002	0.000	0.0001	-0.0006	-0.0008

Table 6. Effects of AI innovation on firm productivity

This table reports results examining the effects of labor-augmenting and labor-displacing AI innovation on firm-level productivity. The dependent variable, *Total Factor Productivity* $_{t,\,t+1}$, is the log-transformed change in TFP from year t to t+1. The independent variables are "type" and "other" patent grants in year t. Columns (1)-(4) report the results from OLS regressions; Columns (5)-(8) report second-stage 2SLS regression results. First-stage 2SLS regression results are reported in Appendix C. Variables are defined as in Table 4. Each regression includes firm-level control variables in the year t-1, patent applications-by-year fixed effects corresponding to the numbers of different "type" and "other" patent applications filed by the firm, and firm fixed effects. All continuous variables are winsorized at 1% and 99%. T-statistics are reported in parentheses. *, **, and *** denote significance at 10%, 5%, and 1%.

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Total Factor Productivity $_{t,\,t+1}$							
	OLS				2SLS (2nd Stage)			
Log # Augmenting	0.028***				0.231*			
	(2.674)				(1.867)			
Log # Non-augmenting	-0.005				-0.051			
	(0.976)				(0.308)			
Log # Scope-augmenting		0.024*				0.309**		
		(1.807)				(2.136)		
Log # Non-scope-augmenting		-0.003				-0.098		
		(0.621)				(0.492)		
Log # Core-augmenting			0.030**				0.229**	
			(2.633)				(2.033)	
Log # Non-core-augmenting			-0.006				0.032	
			(0.968)				(0.139)	
Log # Displacing				0.006				0.034
				(0.340)				(0.272)
Log # Non-displacing				-0.002				0.080
				(0.230)				(0.682)
Size	0.017	0.019	0.017	0.022	0.015	0.015	0.014	0.023**
	(1.371)	(1.488)	(1.351)	(1.589)	(1.508)	(1.577)	(1.411)	(2.357)
MTB	-0.000	-0.000	-0.000	-0.000	-0.000	-0.000	-0.000	0.000
	(0.554)	(0.468)	(0.574)	(0.289)	(0.472)	(1.025)	(0.622)	(0.214)
ROA	-0.067	-0.070	-0.067	-0.072	-0.051	-0.061	-0.051	-0.066*
	(1.597)	(1.638)	(1.593)	(1.639)	(1.305)	(1.411)	(1.238)	(1.831)

R&D	-0.006	-0.006	-0.006	-0.005	-0.004	-0.004	-0.001	-0.009
	(0.542)	(0.558)	(0.534)	(0.489)	(0.481)	(0.549)	(0.142)	(0.922)
R&D missing	-0.023	-0.025	-0.023	-0.024	-0.022	-0.028	-0.017	-0.044
	(0.591)	(0.634)	(0.578)	(0.650)	(0.555)	(0.709)	(0.459)	(1.102)
Observations	13,662	13,662	13,662	13,662	12,964	12,951	12,959	13,017
R-squared	0.279	0.278	0.279	0.278	-0.006	-0.009	-0.006	-0.003
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Kleibergen-Paap F					84.3	30.2	48.8	48.0
# of AI type *k* apps × year FEs	No	No	No	No	Yes	Yes	Yes	Yes
# of other apps × year FEs	No	No	No	No	Yes	Yes	Yes	Yes

Table 7. Labor-displacing AI innovation and cost savings

This table reports results examining the effects of labor-displacing AI innovation on firm-level costs. The dependent variables, *Operating expenses* $_{t,\ t+1}$ and *SG&A* $_{t,\ t+1}$, are the percentage growth rates in total operating expenses and SG&A expenditures, respectively, from year t to t+1. The explanatory variables are "type" and "other" patent grants in year t. Columns (1) and (3) report the results from OLS regressions; Columns (2) and (4) report second-stage 2SLS results. Variables are defined as in Table 4. Each regression includes firm-level control variables in the year t-1, patent applications-by-year fixed effects corresponding to the numbers of different "type" and "other" patent applications filed by the firm, and firm fixed effects. All continuous independent variables are winsorized at 1% and 99%. T-statistics are reported in parentheses. *, **, and *** denote significance at 10%, 5%, and 1%.

	(1)	(2)	(3)	(4)
VARIABLES	Operating expenses $_{t,\ t+1}$		SG&A $_{t,\ t+1}$	
	OLS	2SLS (2nd Stage)	OLS	2SLS (2nd Stage)
Log # Displacing	-0.002	-0.126**	0.003	-0.087**
	(0.530)	(2.485)	(0.941)	(2.126)
Log # Non-displacing	-0.000	0.096**	-0.002	0.051
	(0.090)	(2.261)	(0.732)	(0.975)
Size	-0.047***	-0.048***	-0.037***	-0.041***
	(4.936)	(4.539)	(3.205)	(4.035)
MTB	0.020***	0.019***	0.019***	0.016***
	(5.504)	(4.594)	(6.221)	(8.983)
ROA	0.017	0.021	0.053**	0.052**
	(0.880)	(1.003)	(2.517)	(2.094)
R&D	-0.080***	-0.080***	-0.071***	-0.070***
	(3.142)	(3.117)	(3.624)	(3.801)
R&D missing	-0.258*	-0.258	-0.209**	-0.207**
	(1.711)	(1.659)	(2.343)	(2.279)
Observations	15,693	14,981	13,741	13,044
R-squared	0.223	0.037	0.250	0.047
Firm fixed effects	Yes	Yes	Yes	Yes
Kleibergen-Paap F		58.5		152.3
# of AI type *k* apps × year FEs	No	Yes	No	Yes
# of other apps × year FEs	No	Yes	No	Yes

Table 8. Labor augmentation, labor displacement, and firm value

This table reports results examining the value effects of labor-augmenting and labor-displacing AI innovation. The dependent variable, *Tobin`s Q* $_{t,\ t+1}$, is the net change in Tobin`s Q from year t to t+1. The independent variables are "type" and "other" patent grants in year t. Columns (1)-(4) reports the results from OLS regressions, and Columns (5)-(8) report second-stage results of 2SLS regressions. First-stage 2SLS regression results are reported in Appendix C. Variables are defined as in Table 4. Each regression includes firm-level control variables in year t-1, patent applications-by-year fixed effects corresponding to the numbers of different "type" and "other" patent applications filed by the firm, and firm fixed effects. All continuous variables are winsorized at 1% and 99%. T-statistics are reported in parentheses. *, **, and *** denote significance at 10%, 5%, and 1%.

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Tobin`s Q $_{t,\ t+1}$							
	OLS				2SLS (2nd Stage)			
Log # Augmenting	0.416***				8.151**			
	(3.349)				(2.350)			
Log # Non-augmenting	-0.013				-5.350			
	(0.078)				(0.561)			
Log # Scope-augmenting		0.202**				10.333***		
		(2.015)				(3.053)		
Log # Non-scope-augmenting		0.067				-9.367		
		(0.519)				(1.185)		
Log # Core-augmenting			0.279**				6.507***	
			(2.398)				(3.078)	
Log # Non-core-augmenting			0.048				-13.507	
			(0.327)				(1.603)	
Log # Displacing				0.482***				9.293***
				(3.174)				(4.590)
Log # Non-displacing				0.012				-7.214
				(0.077)				(1.330)
Size	-2.031***	-1.978***	-2.009***	-2.029***	-2.075***	-2.226***	-2.259***	-2.226***
	(3.648)	(3.654)	(3.670)	(3.654)	(3.377)	(3.772)	(4.050)	(3.923)

MTB	-0.158***	-0.157***	-0.157***	-0.157***	-0.153***	-0.153***	-0.151***	-0.153***
	(10.384)	(10.385)	(10.366)	(10.285)	(8.951)	(9.419)	(8.712)	(8.541)
ROA	-4.352**	-4.403**	-4.373**	-4.349**	-3.998*	-3.958**	-3.520*	-3.955*
	(2.176)	(2.187)	(2.189)	(2.176)	(2.000)	(2.034)	(1.759)	(1.894)
R&D	1.194***	1.193***	1.194***	1.193***	1.129**	1.275**	1.188***	1.282**
	(2.707)	(2.723)	(2.715)	(2.699)	(2.403)	(2.563)	(2.668)	(2.602)
R&D missing	2.258**	2.243**	2.257**	2.273**	2.546**	2.981**	2.517**	2.807***
	(2.342)	(2.351)	(2.351)	(2.355)	(2.046)	(2.235)	(2.017)	(2.797)
Observations	15,053	15,053	15,053	15,053	14,299	14,309	14,310	14,334
R-squared	0.125	0.125	0.125	0.125	0.009	0.004	-0.004	0.007
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Kleibergen-Paap F					48.9	64.3	55.9	50.4
# of AI type *k* apps × year FEs	No	No	No	No	Yes	Yes	Yes	Yes
# of other apps × year FEs	No	No	No	No	Yes	Yes	Yes	Yes

Table 9. Labor-augmenting AI innovation, labor market frictions, and firm value

This table reports results examining the effects of augmenting AI innovation on firm value in subsamples. The dependent variable is the net changes in Tobin`s Q from year t to t+1. The independent variables are "type" and "other" patent grants in year t. *Employee Turnover Rate* is defined as the state-level average job separations (scaled by a firm's total employees), excluding the focal firm. *NCC enforceability* is the firm's headquarters state enforceability index for non-compete clauses (NCC), drawn from Garmaise (2011). *Within-occupation transferability* is the industry-level (2-digit SIC) weighted average within-occupation transferability among occupations exposed to AI as defined in Section 4.1. For each occupation, within-occupation transferability is measured as the percentage of job inflow transitions originating from within the 3-digit SOC occupation. Job transition data are from Revelio Labs. Occupational skill requirements are drawn from the O*NET database. Other variables are defined as in Table 4. Each regression includes firm-level control variables in year t-1, patent applications-by-year fixed effects corresponding to the numbers of different "type" and "other" patent applications filed by the firm, and firm fixed effects. All continuous variables are winsorized at 1% and 99%. T-statistics are reported in parentheses. *, **, and *** denote significance at 10%, 5%, and 1%.

Panel A. Labor supply from the external labor market

	(1)	(2)	(3)	(4)	(5)	(6)
	Employee turnover rate					
VARIABLES	High	Low	High	Low	High	Low
Log # Augmenting	14.799***	5.052				
	(3.498)	(1.289)				
Log # Non-augmenting	-6.663	-2.841				
	(1.025)	(0.594)				
Log # Scope-augmenting			10.545***	5.402		
			(2.893)	(1.287)		
Log # Non-scope-augmenting			-10.087	-0.375		
			(0.998)	(0.073)		
Log # Core-augmenting					0.827	1.451
					(0.194)	(0.361)

Log # Non-core-augmenting					-7.358	3.497
					(1.048)	(1.010)
Size	-3.099***	-2.719***	-3.073***	-2.534***	-2.823***	-2.415***
	(4.342)	(3.122)	(3.747)	(3.017)	(4.169)	(2.905)
MTB	-0.153***	-0.175**	-0.150***	-0.173**	-0.149***	-0.179**
	(7.342)	(2.201)	(6.737)	(2.245)	(6.428)	(2.241)
ROA	-2.630	-3.590	-2.624	-4.016	-2.486	-3.848
	(0.756)	(0.889)	(0.715)	(1.011)	(0.710)	(1.010)
R&D	2.148***	0.681**	2.012***	0.519*	2.341***	0.614*
	(2.938)	(2.237)	(2.921)	(1.727)	(3.350)	(1.818)
R&D missing	8.408***	1.064	6.994***	0.572	8.291***	0.404
	(2.836)	(1.555)	(3.274)	(0.758)	(2.825)	(0.441)
Employee turnover rate	1.670	-9.029	0.874	-7.147	0.437	1.463
	(0.675)	(0.759)	(0.418)	(0.741)	(0.236)	(0.129)
Observations	6,054	6,172	6,035	6,175	6,042	6,201
R-squared	0.005	0.008	0.008	0.009	0.014	0.007
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Kleibergen-Paap F	53.2	44.2	72.2	54.9	55.8	32.6
# of AI type *k* apps × year FEs	Yes	Yes	Yes	Yes	Yes	Yes
# of other apps × year FEs	Yes	Yes	Yes	Yes	Yes	Yes

Panel B. External hiring costs

	(1)	(2)	(3)	(4)	(5)	(6)
	NCC enforceability					
VARIABLES	High	Low	High	Low	High	Low
Log # Augmenting	4.588	5.599				
	(1.021)	(1.017)				
Log # Non-augmenting	-12.685	2.155				
	(1.344)	(0.487)				
Log # Scope-augmenting			10.835**	1.415		
			(2.239)	(0.312)		
Log # Non-scope-augmenting			-12.484	-3.930		
			(1.660)	(0.664)		
Log # Core-augmenting					5.227**	-1.701
					(2.132)	(0.400)
Log # Non-core-augmenting					-7.891	-0.440
					(0.977)	(0.119)
Size	-1.449**	-2.105***	-1.471**	-2.254***	-1.258**	-2.119***
	(2.364)	(4.476)	(2.476)	(4.510)	(2.124)	(4.665)
MTB	-0.224***	-0.119***	-0.221***	-0.115***	-0.224***	-0.117***
	(25.071)	(8.353)	(21.991)	(7.241)	(23.155)	(7.770)
ROA	-6.132**	-5.244***	-7.372**	-5.283***	-7.409**	-5.404***
	(2.049)	(3.001)	(2.457)	(3.142)	(2.508)	(3.208)
R&D	1.161*	1.069***	1.241*	1.380***	1.203*	1.208***
	(1.788)	(4.485)	(1.718)	(5.565)	(1.886)	(4.558)
R&D missing	1.226	1.925***	1.651	2.693***	1.207	2.009***
	(0.882)	(3.835)	(1.146)	(3.957)	(0.913)	(3.716)
NCC enforceability	478.366	-2,400.770	-9,012.870	-518.994	-35,246.637	2,497.774
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	6,174	5,935	6,194	5,967	6,199	5,951

R-squared	0.003	0.009	0.005	0.011	0.013	0.010
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Kleibergen-Paap F	38.8	168.8	47.4	90.7	36.5	87.9
# of AI type *k* apps × year FEs	Yes	Yes	Yes	Yes	Yes	Yes
# of other apps × year FEs	Yes	Yes	Yes	Yes	Yes	Yes

Panel C. Within-occupation transferability

	(1)	(2)	(3)	(4)	(5)	(6)
	Within-occupation transferability					
VARIABLES	High	Low	High	Low	High	Low
Log # Augmenting	8.067	3.417				
	(1.648)	(1.391)				
Log # Non-augmenting	9.346**	-21.095*				
	(2.732)	(1.731)				
Log # Scope-Augmenting			5.221	9.967***		
			(1.075)	(3.332)		
Log # Non-Scope-Augmenting			11.362**	-18.151*		
			(2.143)	(1.966)		
Log # Core-Augmenting					11.579*	3.274
					(2.028)	(0.783)
Log # Non-core-augmenting					3.980	-15.564
					(0.765)	(1.380)
Size	-2.431***	-2.435***	-2.328***	-2.587***	-2.522***	-2.187***
	(3.563)	(3.193)	(3.115)	(3.506)	(3.600)	(3.073)
MTB	-2.262***	-0.127***	-2.220***	-0.129***	-2.267***	-0.127***
	(4.022)	(13.663)	(3.888)	(14.993)	(3.978)	(15.447)
ROA	2.699	-4.670	2.353	-5.003	3.019	-4.889
	(1.021)	(1.100)	(0.867)	(1.113)	(1.254)	(1.141)
R&D	1.412**	0.765*	1.624**	1.311***	1.461**	0.605
	(2.085)	(1.866)	(2.507)	(3.404)	(2.077)	(1.352)
R&D missing	4.729**	2.208	5.091***	3.922*	4.557**	2.838
	(2.347)	(0.915)	(2.789)	(1.866)	(2.480)	(1.007)
Within-Occupation transferability	6,313.511***	-2,551.212	6,266.629***	-4,037.233	5,642.760***	-2,245.559
	(3.486)	(0.783)	(3.664)	(1.380)	(2.873)	(0.663)
Observations	6,444	5,933	6,439	5,909	6,434	5,919
R-squared	0.027	-0.013	0.018	-0.006	0.043	-0.000

Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Kleibergen-Paap F	50.6	40.6	66.0	45.1	65.1	38.8
# of AI type *k* apps × year FEs	Yes	Yes	Yes	Yes	Yes	Yes
# of other apps × year FEs	Yes	Yes	Yes	Yes	Yes	Yes

Table 10. Labor-displacing AI innovation, labor-market frictions, and firm value

This table reports results examining the effects of displacing AI innovation on firm value in subsamples. The dependent variable is the net change in Tobin`s Q from year t to t+1. The independent variables are "type" and "other" patent grants in year t. *NCC enforceability* is the enforceability index (from Garmaise, 2011) of non-compete clauses (NCC) in the firm's headquarters state. *Unemployment insurance* is calculated as the product of average weekly benefits and recipiency rates of regular UI programs. State-level unemployment insurance data are obtained from the U.S. Department of Labor (https://oui.doleta.gov/unemploy/data.asp). Other variables are defined as in Table 4. Each regression includes firm-level control variables in year t-1, patent applications-by-year fixed effects corresponding to the numbers of different "type" and "other" patent applications filed by the firm, and firm fixed effects. All continuous variables are winsorized at 1% and 99%. T-statistics are reported in parentheses. *, **, and *** denote significance at 10%, 5%, and 1%.

	(1)	(2)	(3)	(4)
	Unemployment insurance		NCC enforceability	
VARIABLES	High	Low	High	Low
Log # Displacing	-3.486	11.332*	4.693*	14.218
	(0.714)	(1.766)	(1.691)	(1.540)
Log # Non-displacing	-12.706	-14.425*	-3.919	-3.434
	(1.488)	(1.926)	(0.883)	(0.407)
Size	-1.901**	-2.380**	-1.375**	-2.334***
	(2.215)	(2.074)	(2.375)	(3.730)
MTB	-1.021***	-0.150***	-0.219***	-0.221
	(8.681)	(28.979)	(17.811)	(1.278)
ROA	-3.914**	-7.957***	-7.141**	-5.672***
	(2.253)	(2.703)	(2.015)	(2.928)
R&D	0.772	2.361*	1.120*	1.109
	(1.127)	(1.945)	(1.900)	(1.482)
R&D missing	3.162	2.823	2.222*	3.259
	(1.477)	(0.961)	(1.798)	(1.459)
Unemployment insurance	0.000***	-0.001***		
	(2.727)	(4.700)		
NCC enforceability index			-3,001.920	-1,058.947
			(0.000)	(0.000)
Observations	4,938	4,848	5,305	5,106
R-squared	-0.006	0.014	0.019	0.005
Firm fixed effects	Yes	Yes	Yes	Yes
Kleibergen-Paap F	34.4	30.5	38.1	63.3
# of AI type *k* apps × year FEs	Yes	Yes	Yes	Yes
# of other apps × year FEs	Yes	Yes	Yes	Yes

Table 11. Labor-augmenting AI innovation, skill characteristics, and firm value

This table reports results examining the effects of Augmenting AI innovation on firm value in subsamples. The dependent variable is the net changes in Tobin`s Q from year t to t+1. The independent variables are the (log-transformed) numbers of "type" and "other" patent grants in year t. In Panel A, Columns (1), (3), and (5) report the results for the subsample of firms that are either headquartered in a state with high skill proximity or operate in an (2-digit SIC) industry with high skill proximity. Columns (2), (4), and (6) report results for the subsample of firms that are located in a state with low skill proximity and operate in a (2-digit SIC) industry with low skill proximity. Skill proximity for a pair of occupations is equal to one minus skill distance, which is calculated as follows. Occupational-pair-level skill distance is the number of unique skills (weighted by skill importance from the O*NET database) in the occupation relative to all other pair occupations. *Local (Industry) skill distance* is the average (weighted by occupational employment share) skill distance across occupations exposed to "type" AI innovations. In Panel B, Columns (1), (3), and (5) report the results for the subsample of firms that have either a high level of internal skill supply or operate in a (2-digit SIC) industry with a high level of skill supply. Columns (2), (4), and (6) report the results for the subsample of firms that have a low level of internal skill supply and operate in a (2-digit SIC) industry with a low level of skill supply. Occupation-level skill supply is calculated as the number of important skills (importance >= 4) in the occupation. *Local (Industry) Skill supply* is the average (weighted by occupational employment share) skill supply across occupations that are exposed to "type" AI innovations. Other variables are defined as in Table 4. Each regression includes firm-level control variables in year t-1, patent applications-by-year fixed effects corresponding to the numbers of different "type" and "other" patent applications filed by the firm, and firm fixed effects. Continuous variables are winsorized at 1% and 99%. T-statistics are given in parentheses. *, **, and *** denote significance at 10%, 5%, and 1%.

Panel A. Skill proximity

	(1)	(2)	(3)	(4)	(5)	(6)
	Skill proximity					
VARIABLES	High	Low	High	Low	High	Low
Log # Augmenting	14.125***	8.289**				
	(3.463)	(2.260)				
Log # Non-augmenting	-17.745	8.507				
	(1.236)	(1.309)				
Log # Scope-augmenting			8.994***	2.530		
			(2.858)	(0.215)		
Log # Non-scope-augmenting			-19.354*	4.238		
			(1.717)	(1.054)		
Log # Core-augmenting					8.604**	1.364
					(2.214)	(0.167)

Log # Non-core-augmenting					-14.981	13.147**
					(1.546)	(2.520)
Size	-1.597**	-3.884***	-1.660**	-3.985***	-1.580*	-3.803***
	(2.178)	(3.868)	(2.206)	(6.553)	(1.913)	(4.223)
MTB	-0.214***	-2.008***	-0.214***	-1.912***	-0.214***	-2.005***
	(19.779)	(5.548)	(17.537)	(5.778)	(18.814)	(6.625)
ROA	-6.026*	-2.474	-6.011*	-2.890	-6.056*	-2.887
	(1.745)	(0.556)	(1.753)	(0.672)	(1.851)	(0.638)
R&D	0.578	3.536**	0.714**	3.982***	0.782**	3.859**
	(1.624)	(2.305)	(2.360)	(3.186)	(2.225)	(2.580)
R&D missing	1.912	7.699***	2.271*	8.251***	1.901	8.647***
	(1.329)	(3.015)	(1.749)	(3.777)	(1.459)	(3.502)
Local skill specificity	4.866	-4.072	6.218	-5.603*	4.603	-3.303
	(0.850)	(0.846)	(1.174)	(1.878)	(1.165)	(0.771)
Industry skill specificity	4.032	-1.720	4.109	0.241	3.709	-2.673
	(0.661)	(0.776)	(0.665)	(0.091)	(0.613)	(1.176)
Observations	7,387	2,739	7,353	2,741	7,382	2,746
R-squared	-0.014	0.019	-0.015	0.041	-0.004	-0.010
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Kleibergen-Paap F	33.5	12.4	37.0	30.8	38.3	20.5
# of AI type *k* apps × year FEs	Yes	Yes	Yes	Yes	Yes	Yes
# of other apps × year FEs	Yes	Yes	Yes	Yes	Yes	Yes

Panel B. Skill supply (from internal and external labor markets)

	(1)	(2)	(3)	(4)	(5)	(6)
	Skill supply					
VARIABLES	High	Low	High	Low	High	Low
Log # Augmenting	5.920***	5.745				
	(3.324)	(0.889)				
Log # Non-augmenting	-7.962	-1.368				
	(0.585)	(0.547)				
Log # Scope-augmenting			6.544*	-0.981		
			(1.834)	(0.414)		
Log # Non-scope-augmenting			-18.914	0.543		
			(1.286)	(0.274)		
Log # Core-augmenting					9.805***	3.546
					(3.039)	(0.554)
Log # Non-core-augmenting					-15.343	-1.766
					(1.326)	(0.705)
Size	-2.473***	-0.559	-2.842***	-0.550	-2.731***	-0.516
	(3.182)	(0.366)	(4.146)	(0.365)	(3.450)	(0.353)
MTB	-0.216***	-1.510***	-0.221***	-1.547***	-0.222***	-1.508***
	(7.648)	(5.610)	(8.079)	(5.619)	(8.593)	(5.723)
ROA	-5.873*	-0.028	-5.767	-0.405	-4.957	-0.442
	(1.785)	(0.018)	(1.656)	(0.240)	(1.534)	(0.294)
R&D	1.341**	0.492	1.764***	0.320	1.583**	0.505
	(2.217)	(0.990)	(3.328)	(0.618)	(2.108)	(1.094)
R&D missing	2.370	0.828	3.581	0.826	3.120	0.638
	(1.373)	(0.699)	(1.497)	(0.697)	(1.239)	(0.507)
Skill supply within industry	74.466	-34.934	60.341	-22.083	63.745	-34.918
	(0.831)	(0.744)	(0.909)	(0.368)	(0.799)	(0.849)
Skill supply within firm	-0.458	0.351	-0.575	0.406	-0.396	0.387
	(0.576)	(1.511)	(0.709)	(1.411)	(0.533)	(1.585)

Observations	7,431	2,548	7,464	2,547	7,451	2,554
R-squared	0.012	0.029	-0.007	0.032	-0.001	0.032
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Kleibergen-Paap F	26.7	46.8	41.2	42.8	45.2	41.1
# of AI type *k* apps × year FEs	Yes	Yes	Yes	Yes	Yes	Yes
# of other apps × year FEs	Yes	Yes	Yes	Yes	Yes	Yes