

Brief for the Canada House of Commons Study on the Implications of Artificial
Intelligence Technologies for the Canadian Labor Force
**Generative Artificial Intelligence Shatters Models of AI
and Labor**

Morgan R. Frank^{1,2,3,*}

¹*Department of Informatics and Networked Systems, University of Pittsburgh, Pittsburgh, PA 15216 USA*

²*Digital Economy Lab, Stanford University, Stanford, CA 94305 USA*

³*Connection Science, Massachusetts Institute of Technology, Cambridge, MA, USA*

**To whom correspondence should be addressed. E-mail: mrfrank@pitt.edu*

November 2023

1 Introduction

Exciting advances in generative artificial intelligence (AI) have sparked concern for jobs, education, productivity [1], and the future of work. As with past technologies, generative AI may not lead to mass unemployment. But, unlike past technologies, generative AI is creative, cognitive, and potentially ubiquitous which makes the usual assumptions of automation predictions ill-suited for today. Existing projections suggest that generative AI will impact workers in occupations that were previously considered immune to automation. As AI's full set of capabilities and applications emerge, policy makers should promote workers' career adaptability. This goal requires improved data on job separations and unemployment by locality and job titles in order to identify early-indicators for the workers facing labor disruption. Further, prudent policy should incentivize education programs to accommodate learning with AI as a tool while preparing students for the demands of the future of work.

2 Generative AI Shatters Conventional Wisdom

Rarely is a single technology so pervasive as to make whole occupations obsolete. Rather, technology performs specific tasks and impacts the demand for workers who perform those tasks in their job (i.e., skill-biased technological change [2, 3, 4]). Thus efforts to predict technology-driven job disruptions focus on the variation tasks and activities across occupations [5, 6]. However, knowing both the specific capabilities of technology and the activities of workers across industries, firms,

and regions has been a difficult challenge. Historically, researchers have relied on empirical heuristics that broadly characterize the relationship between classes of workers and the technologies they interact with. Routine physical workers (e.g., in manufacturing and construction) are thought to be replaced by robotics, while non-routine cognitive workers become more innovative and productive with computing and machine learning.

These heuristics inform predictions of workers' exposure to technology. For example, creative workers—including artists, graphic designers, and even scientists—are typical examples of non-routine cognitive occupations that have been assumed to be safe from automation. First appearing as a pre-print in 2013, a high-profile study [7] from Oxford University asserted:

“Because creativity, by definition, involves not only novelty but value, and because values are highly variable, it follows that many arguments about creativity are rooted in disagreements about value. . . In the absence of engineering solutions to overcome this problem, it seems unlikely that occupations requiring a high degree of creative intelligence will be automated in the next decades.”

Rather than deviating from these norms, automation studies differ mostly in their modeling of technology's capabilities (e.g., based on expert opinion [8], public opinion [9], or patents [10, 11]).

Do these heuristics hold with today's technology? Or does generative AI shatter conventional wisdom? While we are still waiting to see the full capabilities and applications of generative AI, large language models (LLMs), such as OpenAI's ChatGPT and Google's Bard, and image generators, such as Midjourney and OpenAI's DALL-E, can draft essays, write computer code, and perform graphic design. These capabilities are clear examples of what would traditionally be called non-routine cognitive work and suggest that a previously immune class of white collar workers may now compete with automation. For example, a recent study [12] from OpenAI and the University of Pennsylvania finds that highly-educated U.S. workers in high-earning occupations are most exposed to LLMs (i.e., these occupations require workplace activities that could be performed by LLMs).

Generative AI shatters the conventional wisdom of past automation studies. For instance, exactly one decade after the quote above, AI is performing creative work in the form of writing, image creation, and software development [13]. AI has evolved and so too must our approach to economics research, education programs, and labor policy. Rather than fallible heuristics, a response to AI must be informed by detailed empirics representing workers across the economy.

3 Improve Data on Negative Labor Outcomes

New applications of generative AI increase the need for labor data that liberates policy from limiting assumptions regarding automation. Currently, automation studies compare occupation exposure scores to decreasing employment or wages [7, 14, 8, 9, 10, 15], but this approach merely reflects the data that has been available rather than the direct symptoms of AI exposure. For example, unemployment or job losses associated with an occupation can increase or decrease even as employment for that occupation increases [16]. Thus, economic policy must be informed by the direct symptoms of workers' exposure to AI.

Generative AI can do certain creative and cognitive activities which may diminish the demand for workers to perform those activities. Rather than automating an entire occupation, a more




Figure 1: The pipeline of technologies’ most to least common impact on workers. Technology reshapes the demand for specific skills rather than whole occupations. If workers fail to adapt their skills, then job separations can occur leading workers to seek new jobs. Workers may apply for an receive unemployment benefits if new jobs are not found. Data is needed to describe each part of this impact pipeline for each occupation.

direct outcome is that workers shift their activities to complement the work of AI [6] (i.e., within-occupation skill change). If workers cannot adapt, then job separations (i.e., employer firing or worker quitting) may displace workers [17, 18, 19]. Displaced workers who cannot quickly find new employment may rely on unemployment benefits while they continue job seeking. Figure 1 describes this pipeline of AI impact from the most direct and common symptoms of AI exposure to the least common. Note, however, that this pipeline says nothing about employment or wages.

For most national economies, data for describing the pipeline of AI impact does not exist at the level of detail needed to improve automation predictions. But there are steps to improve the level insights that can inform labor policy in response to generative AI.

Improved data on within-occupation skill change must be dynamic enough to model the new skills and capabilities that AI will enable (e.g., prompt engineering: the skill of writing prompts for LLMs) while also subject to regular updates and maintaining representativeness of workplace activities across the economy. For example, the U.S. Bureau of Labor Statistics (BLS) reports skill and task requirements for over 900 different occupations in the O*NET database; this database achieves good representation for workers across U.S. sectors but each occupation is only updated every five years and the taxonomy of workplace skills and tasks is updated even less frequently. With this data alone, it would be difficult for policy makers to respond to emergent demand for new workplace activities, such as “prompt engineering.” Researchers have begun to use job postings to supplement federal data, but, currently, job postings can be biased towards certain industries and offer their own data challenges (e.g., processing and cleaning free text job descriptions).

Improved data on job separations requires real-time information on the employers or occupations where job separations are occurring, the nature of the separation (i.e., worker quitting their job or employer firing a worker), and reason for the separation (e.g., worker was not adaptable to new workplace demands versus a human resource issue). For example, the U.S. BLS Job Openings and Labor Turnover Survey (JOLTS) program produces data on job openings, hires, and separations by month and state, but does not detail the industry, employer, or occupation related to these important labor dynamics. Canada’s provincial economies are diversified enough to require an explanation of job separations in order to better tie job losses to generative AI compared to other sources of disruption. Online worker resumes (e.g., from LinkedIn or Indeed.com) offer a potential data source with real-time insights into job separations, separated workers’ skills, and efforts to seek new employment, but, as with job postings, representativeness and bias towards certain industry may pose a challenge.

Improved data on unemployment should go beyond total unemployment by labor market to also include *unemployment risk* by occupation or industry. Traditionally, “unemployment by occupation” is a strange concept because, by definition, unemployed individuals have no occupation. However, the probability of receiving unemployment by occupation should suffice and, again, enable policy makers and economists to directly tie AI disruptions to unemployment versus other sources. Typically, unemployment insurance benefits are contingent on recipients’ continued effort to seek new employment and, in service to this goal of re-employment, insurance offices often track the employment history of benefit claimants. Combined with total unemployment statistics and employment shares by occupation, this data could create new insights into the causes of unemployment and portions of the workforce that are in greatest need of policy intervention with the increased adoption of generative AI.

4 Adapting Education

Generative AI will reshape the skill requirements of white collar occupations and so the educational pipeline must adapt if today’s students are to fill these roles in the future. As professors and teachers struggle to detect LLM-generated essays in classrooms, a more practical approach should incorporate generative AI applications as learning tools, much like calculators in a mathematics. It is not that students should completely depend on generative AI, but learning the tools of the day should be part of any curriculum that prepares students for the future of work.

While workforce statistics abound, comparatively fewer economy-wide data reflect the skills taught during higher education and their influence on long-term career outcomes. For example, the U.S. Department of Education reports enrollment demographics and two-years post-graduation earnings by year, university/college, and major, but offers little insight into career outcomes thereafter or the differences in skills taught within the same major across different universities. Just as researchers compare technology patents occupations’ skills [10, 11], course descriptions from syllabi may offer a scalable option to infer taught skills from text course descriptions. In turn, the methods for studying technology in the labor market should be applied to taught skills to identify the majors and universities with the greatest exposure to generative AI. These educational pathways will require the most attention as we continue to learn about generative AI’s capabilities and applications. Teachers may focus less on skills that are completely automated by AI. But, more commonly, skills related to the capabilities of AI will require more attention so that students can complement technology once they enter the workforce. For example, recent case studies [20, 21] suggest that generative AI tools do not improve the performance of experts but, instead, raise the performance of non-experts to comparable levels. If this observation holds across contexts, then incorporating generative AI into learning curricula has the potential to improve learning objectives for under-performing students [22, 23].

Just as policy makers and government agencies measure skills and labor market outcomes, they should equally measure the sources of labor market skills. In that case of white collared workers who are most exposed to generative AI [12], many of these skills result from a college education. Quantifying the pipeline for specific skills will inform policy to both re-skill today’s workers and to prepare today’s students for the future of work with generative AI.

References

- [1] Martin Neil Baily, Erik Brynjolfsson, and Anton Korinek. Machines of mind: The case for an AI-powered productivity boom. 2023.
- [2] John Hicks. *The theory of wages*. Springer, 1963.
- [3] Giovanni L Violante. Skill-biased technical change. *The new Palgrave dictionary of economics*, 2:1–6, 2008.
- [4] Daron Acemoglu and David Autor. Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics*, volume 4, pages 1043–1171. Elsevier, 2011.
- [5] Ahmad Alabdulkareem*, Morgan R. Frank*, Lijun Sun, Bedoor AlShebli, César Hidalgo, and Iyad Rahwan. Unpacking the polarization of workplace skills. *Science Advances*, 4(7), 2018. (*authors contributed equally).
- [6] Morgan R Frank, David Autor, James E Bessen, Erik Brynjolfsson, Manuel Cebrian, David J Deming, Maryann Feldman, Matthew Groh, José Lobo, Esteban Moro, et al. Toward understanding the impact of artificial intelligence on labor. *Proceedings of the National Academy of Sciences*, 116(14):6531–6539, 2019.
- [7] Carl Benedikt Frey and Michael A. Osborne. The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114:254 – 280, 2017.
- [8] Erik Brynjolfsson, Tom Mitchell, and Daniel Rock. What can machines learn, and what does it mean for occupations and the economy? In *AEA Papers and Proceedings*, volume 108, pages 43–47, 2018.
- [9] Edward W Felten, Manav Raj, and Robert Seamans. A Method to Link Advances in Artificial Intelligence to Occupational Abilities. In *AEA Papers and Proceedings*, volume 108, pages 54–57, 2018.
- [10] Michael Webb. The impact of artificial intelligence on the labor market. *Available at SSRN 3482150*, 2019.
- [11] Benjamin Meindl, Morgan R Frank, and Joana Mendonça. Exposure of occupations to technologies of the fourth industrial revolution. *arXiv preprint arXiv:2110.13317*, 2021.
- [12] Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. GPTs are GPTs: An early look at the labor market impact potential of large language models. *arXiv preprint arXiv:2303.10130*, 2023.
- [13] Ziv Epstein, Aaron Hertzmann, Investigators of Human Creativity, Memo Akten, Hany Farid, Jessica Fjeld, Morgan R Frank, Matthew Groh, Laura Herman, Neil Leach, et al. Art and the science of generative AI. *Science*, 380(6650):1110–1111, 2023.

- [14] Melanie Arntz, Terry Gregory, and Ulrich Zierahn. The Risk of Automation for Jobs in OECD Countries. Technical report, 2016.
- [15] Songül Tolan, Annarosa Pesole, Fernando Martínez-Plumed, Enrique Fernández-Macías, José Hernández-Orallo, and Emilia Gómez. Measuring the occupational impact of AI: tasks, cognitive abilities and AI benchmarks. *Journal of Artificial Intelligence Research*, 71:191–236, 2021.
- [16] Monica Mean and Marilia G Mochel. Job automation risk and the future of skills: Skills and competency change in the US workforce. Technical report, US Bureau of Labor Statistics, 2023.
- [17] Dale T Mortensen and Christopher A Pissarides. Technological progress, job creation, and job destruction. *Review of Economic dynamics*, 1(4):733–753, 1998.
- [18] Boyan Jovanovic. Job matching and the theory of turnover. *Journal of political economy*, 87(5, Part 1):972–990, 1979.
- [19] Frank M Fossen and Alina Sorgner. New digital technologies and heterogeneous wage and employment dynamics in the United States: Evidence from individual-level data. *Technological Forecasting and Social Change*, 175:121381, 2022.
- [20] Fabrizio Dell’Acqua, Edward McFowland, Ethan R Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Kraymer, François Candelon, and Karim R Lakhani. Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. *Harvard Business School Technology & Operations Mgt. Unit Working Paper*, (24-013), 2023.
- [21] Erik Brynjolfsson, Danielle Li, and Lindsey R Raymond. Generative AI at work. Technical report, National Bureau of Economic Research, 2023.
- [22] Ethan Mollick and Lilach Mollick. Assigning AI: Seven Approaches for Students, with Prompts. *arXiv preprint arXiv:2306.10052*, 2023.
- [23] Ethan R Mollick and Lilach Mollick. Using AI to implement effective teaching strategies in classrooms: Five strategies, including prompts. *Including Prompts (March 17, 2023)*, 2023.