

Addressing Deskilling as a Result of Human-AI Augmentation in the Workplace

Firuza Huseynova¹

¹ University of Western Ontario, 1151 Richmond St., London, N6A 3K7, Canada

Abstract

The integration of artificial intelligence (AI) technology into the workplace has become a focal point of discussion in recent years. While management scholars have typically advocated for an approach that augments rather than substitutes human labour, little consideration has been given to the potential drawbacks of human-AI augmentation strategies. Thus, this paper focuses on addressing employee deskilling that can arise from human-AI augmentation in the workplace. Drawing on insights from eight semi-structured interviews with experts across industry and academia, this paper details two key factors influencing managerial decision-making in augmentation projects and touches on the blurred lines between augmentation and substitution strategies—especially when augmentation erodes the satisfaction of human work.

Keywords

Artificial intelligence, augmentation, automation, technological unemployment, deskilling, AI ethics, AI governance

1. Introduction

Recent advances in artificial intelligence (AI) technology have led to a greater capacity for machines to replicate nonroutine and nonrepetitive tasks which were once considered impossible to automate [2, 9]. This has sparked significant debate on the role of automated systems in the workplace, particularly around issues of technological unemployment, deskilling, and task encroachment [3, 4]. When it comes to the integration of AI into the workplace, management scholars typically advocate for an approach that complements and augments human labour rather than substitutes it [1]. However, the increasing potential for AI to augment humans in more and more types of work has made it important to study human-AI augmentation from a technology ethics perspective. Thus, my research question is as follows: *How does deskilling arise in the workplace when AI is deployed to augment rather than substitute human labour?*

This paper will begin with an overview of the broader debate between augmentation and substitution, leading to a discussion on deskilling in the workplace. I will then outline the methodology used to conduct eight semi-structured interviews with industry professionals and AI researchers, setting the stage for a discussion of key findings. These key findings center around the types of deskilling seen in the workplace, the key factors influencing managerial decision-making around automation, and the blurred lines between augmentation and substitution. I will close the paper with two recommendations for employers, synthesized from interview insights, on mitigating potential negative consequences of deskilling in the workplace.

2. Background

2.1. The Augmentation vs. Substitution Debate

Among management scholars, the relationship between augmentation and substitution is usually described as a trade-off [3]. Automating a task (used interchangeably with ‘substituting’ for the purposes of this paper) involves handing it over to a machine with little to no human involvement, usually for the sake of more efficient or productive operations [1]. On the other hand, augmenting a

7th Conference on Technology Ethics (TETHICS2024), November 6–7, 2024, Tampere, Finland

✉ fhusecno@uwo.ca



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

task implies introducing heavier human-machine involvement, leading to an ostensibly complementary balance between human and machine capabilities.

In a review of three New York Times best-selling books on the intersection of AI and business, Raisch and Krakowski found that management scholars typically advocate for an augmentation approach rather than resorting to substitution in the form of mass layoffs [1]. In public communications, spokespeople for tech companies such as Microsoft and Google frequently state that companies will use AI in the workforce to complement human abilities rather than replace or restrict them [1]. IBM CEO Ginni Rometty suggested using the term “augmented intelligence” instead of “artificial intelligence”, driving home the point that AI implementation should be aimed at augmenting the existing capabilities of humans rather than replacing them [1].

Additionally, according to a participatory study conducted with 54 knowledge workers across seven fields, workers tend to hold favourable views on human-in-the-loop (HITL) approaches, which aim to ensure an element of human oversight is integrated within workplace augmentation strategies [10]. Proponents of the augmentation view generally argue that humans and AI should coexist in the workplace; if humans do not learn to work with AI, they may risk falling behind the skill curve.

2.2. Deskilling

Although the impact of emerging technologies on human skillsets is very ambiguous and difficult to quantify, potential deskilling as a result of human-AI augmentation should concern ethicists. The concept of deskilling was formed in the 20th century to explain how automation leads to a loss of practical knowledge and artisan skill sets such as looming or tilemaking [5]. Outside of technical skills, Shannon Vallor’s research in moral deskilling suggests that moral skills are just as vulnerable to disruption by technological advancements as other types of skills [5]. On an aggregate level, deskilling is linked with technological unemployment, which, in its structural form, has the potential to disrupt individual well-being, social cohesion, and entire economic systems [4]. The negative consequences of deskilling from both an individual and societal level make it crucial to consider human-AI augmentation in the workplace a legitimate subject matter to study for technology ethicists.

Machines are increasingly encroaching into tasks which require complex cognitive and emotional capacities, previously only possible by humans. Task encroachment refers to the widening of tasks that machines will be able to perform, from mere manual tasks to more cognitive and emotional ones. [2]. Researchers typically distinguish between three types of AI systems: mechanically-intelligent (manual), thinking-intelligent (cognitive) and feeling-intelligent (emotional) [6]. Manual tasks include inspecting and operating equipment, controlling machines, scheduling work or activities, handling or moving objects, recording information, and performing administrative activities [6]. Cognitive tasks include analyzing data or information, processing information, developing objectives or strategies, thinking creatively, interpreting the meaning of information for others, solving problems, and providing consultation or advice [6]. Lastly, emotional work involves communicating with peers or clients, assisting and caring for others, resolving conflicts, negotiating with others, developing or building teams, training or teaching other people, and establishing interpersonal relationships [6]. Today, cognitive and emotional tasks which were once considered impossible to automate have been automated, leading to deskilling in areas like customer service and nursing [5].

Due, in part, to task encroachment, workers in many industries are facing a redundancy of core skills they trained for years to develop [5, 10]. This is partly because businesses focus on automating away tasks to achieve cost efficiencies in the short run, which signals industry competitors to pursue substitution (ie. lay workers off) in order to remain cost-competitive and secure their position in the market [10]. However, in the long run, employees may lose some of the skills required to oversee and alter the processes of AI-enabled tools—or, lose their jobs altogether.

The term ‘technological unemployment’, popularized by John Maynard Keynes in the early 20th century, can be divided into its frictional and structural forms [2, 12]. Frictional technological unemployment occurs when workers shift between jobs in the pursuit of finding their ideal occupation [2]. Short-term frictional unemployment is considered natural and healthy in a market-based economy [2]. However, structural technological unemployment is a more long-term form of unemployment resulting from a mismatch between the skills workers possess and the skills

employers are looking for [2]. Unlike previous eras in history where new technologies eventually created more jobs than they displaced, the current pace of technological advancement, particularly in AI and machine learning, raises concerns about the ability of humans to upskill quickly enough [4]. The potential impact of structural technological unemployment extends beyond economic measures—it can exacerbate inequality, strain social safety nets, and erode the dignity of work [4].

While deskilling is typically a term applied to workers who have lost their jobs due to automation, it can also apply to workers who have experienced their work be augmented by AI. In this paper, I argue that it is misleading to present the difference between augmentation and substitution as a clean-cut trade-off. The lines between augmentation and substitution are blurred, and deskilling can also occur in workers augmented by AI, as will be discussed further on.

3. Methodology

3.1. Research Strategy

This paper relies on an inductive approach, wherein theories and interpretation are the outcome of theory [7]. As to be expected when conducting inductive analysis, the central research question changed multiple times throughout its journey. There was no central theory guiding the development of research, but the research process was approached from a theory-building perspective [7]. Thus, data was gathered and examined first and only then were theories crafted from observations.

3.2. Literature Search

The literature search was primarily collected through two databases: Google Scholar and Western University’s OMNI database. Thanks to the advanced search engine capabilities of both databases, a thorough search was done on combinations of keywords including: “AI”, “deskilling”, “automation”, “augmentation”, “substitution”, and “task encroachment”.

Since the field of information systems is highly interdisciplinary, it was essential to search for studies outside the field to uncover as much relevant research as possible, including fields like business ethics, sociology, and human-computer interaction. To obtain high-quality literature, results were filtered on peer-reviewed academic materials and journals. As well, given the high-paced growth of new technologies, priority was given to articles published in the last five years. Lastly, publications with a higher number of citations were prioritized, as this metric was taken as a proxy for how credible and influential an article was in its respective field.

3.3. Interviews

To add depth to research findings, semi-structured interviews were conducted with two target groups: managers with experience in AI-implementation projects, and researchers focused on AI ethics. A semi-structured interview is one in which the interviewer is able and willing to switch the order of questions asked or ask unique questions to each interviewee depending on their experiences [7]. As this project’s inductive investigation was based on asking open questions, semi-structured interviews were chosen to allow for flexibility and adaptability during the interview process.

A total of eight interviews were conducted. Interviewees were primarily selected based on their experience in the field. See Table 1 for a complete list of interviewees along with their credentials (names eliminated for anonymity).

Table 1
Interviewee Overview

Number	Experiences
I1	Chief AI Officer at major Canadian university
I2	CEO of soft skills training company, doctoral candidate researching AI and society
I3	Independent journalist, host of award-winning tech podcast
I4	Research and policy analyst at Responsible AI Institute (RAI)
I5	AI strategist at Big 5 Canadian bank, alumnus of Oxford Internet Institute

I6	Doctoral candidate at McGill University researching AI governance
I7	Doctoral candidate at University of Toronto researching AI ethics and policy
I8	Senior digital workforce transformation consultant at Big 4 audit firm

3.4. Interview Protocol

Before commencing each interview, consent was established to record audio and take notes during the interview. Interviewees were informed that the collected material would be processed anonymously and confidentially to safeguard identities. The interview protocol contained six base questions, although the questions were adjusted liberally to align with the interviewee's relevant experiences and the overall direction of the discussion. The questions are listed below:

1. What is your name, what does your organization do, and what is your role in the organization?
2. Are you involved in any AI implementation projects? If so, which projects are you involved in?
3. What do you think are the most important factors a business leader should consider when making decisions about automating tasks currently performed by humans?
4. An alternative to substituting labour is using technology to augment people's work. But, even when work is augmented, skills can be eroded. How do you decide which skills people need to keep and which ones are okay to lose?
5. Which type of workers are more likely to lose skills when AI is introduced to augment their workforce?
6. What are new skills that workers or employees need to develop to adapt to the adoption of AI in the organization?

3.5. Limitations

When it comes to selecting a research topic, one's decision about what to investigate always precedes the methodological issues of how the research might best continue [7]. It is undoubtable that many researchers hold values driving them to choose their research topic, especially in the case of a topic related to ethics, as mine is. However, I do not subscribe to the view that research should be values-free; on the other hand, I believe value-commitment is a good thing for researchers to have as they can use their passions to direct and interpret their investigations while remaining open about the types of biases they may have. In my free time, I mainly consume content (ie. books, podcasts, and videos) that is critical of Big Tech and technological determinism. Thus, my paper reflects this perspective. However, I believe I have developed a fair and balanced paper that considers the perspectives of techno-optimists and tech skeptics alike.

Still, one limitation of the paper is the semi-structured nature of the interviews led to a lack of standardization among many interview questions. Although valuable insights were gained from each individual interview, it was difficult to compare and contrast answers given to any of the questions in a statistically significant way. However, the qualitative nature of the data allowed for a deeper exploration of diverse perspectives, ultimately providing a nuanced understanding of the impact of AI on the workforce.

4. Findings

4.1. Context-Dependency

I2: "In scholarship, the goal is new knowledge creation. I'm not going to copy and paste exactly what GPT gives me into a paper because that is not new knowledge. And that's not the goal of academia. However, I will copy & paste at [soft skill development startup] because that is not the goal. The goal is that [clients] walk away having learned a skill. And if GPT can make that easier for us, sign me up."

While there were significant differences in the way each interviewee approached the questions, a common theme was the importance of context-dependency when deciding which tasks to automate away. Specifically, interviewees 1 and 4 (hereby referred to as I1 and I4) mentioned that a one-size-fits-all approach would not be optimal for dynamic work environments, meaning automation strategies would need to be tailored to suit the specific requirements of different departments and teams. I2 also noted that the goals of a particular sector (private, public, non-profit, or academia) could help determine whether AI should be used to augment tasks like writing and content generation. Moreover, I8, a senior Digital Workforce Transformation consultant, mentioned that the size of a company is a key factor in their willingness to adopt human-AI augmenting technologies in the first place. They stated that a small startup with a leader who is tapped into the tech space will be much more likely to opt into human-AI augmentation in the workforce, as opposed to a company that has been operating with engrained processes for 50+ years, which may be more reluctant or risk-averse to testing new technologies in the workplace. Overall, the impact of automation can vary significantly based on the unique environment, industry, workforce composition, and strategic objectives of a given company.

4.2. Types of Deskilling

I5: “This is like the futurist’s dream, right? The machines take care of all the automated, mindless work. And then we as humans can be creative and focus on other things. Will that happen in actuality? I don’t know. It’s not a very black or white thing. The answer will lie somewhere in between.”

Firstly, most responses indicated that deskilling in areas of manual labour and mechanical data-entry is largely positive since it frees humans to focus on more meaningful cognitive and emotional work. Specifically, I5 stated that skills in calculation or summarization will likely be devalued, but this will only lead to a greater capacity for humans to focus on honing in their analytical and problem-solving skills. This belief is consistent with the optimistic view of human-AI augmentation on the labour market, being that it will create new jobs as the demand for nonroutine and nonrepetitive skills increases [6]. Thus, labour has the potential to ‘upgrade’ from being mechanical to cognitive.

However, when it comes to the more creative manual skills that may be lost in this transition (ie. artisanal human craftwork), I2 argued that there will always be a market for authenticated time-consuming human work. They bring up the example of jewelry: buying a ring from someone who spent 50 hours forging and faceting a gem is completely different than buying something stamped out in a factory. They believe that there will always be demand for human-crafted language, code, and prose as we value things just for the sake of being made by humans.

Lastly, I8 noted the importance of the quality of AI tools available in determining how deskilled a workforce becomes. If workers have high-quality AI tools that can summarize information or obtain insights more effectively, they are more likely to become deskilled at completing those tasks. This is simply because if a worker does not have to read an entire document to understand complex terminology or summarize it, then they are not practicing reading comprehension and summarization—essentially becoming deskilled in those tasks.

4.3. Key Factors Influencing Managerial Decision-Making

4.3.1. Time and Effort

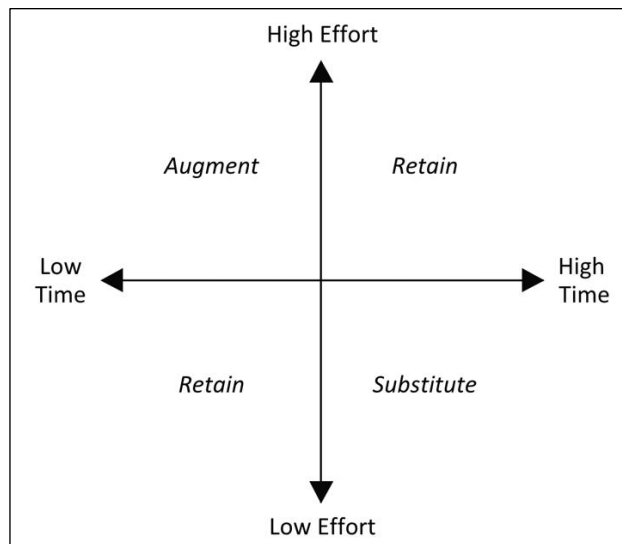
I1: “Use the technology to skill up the people who are in place and say — look, save your brain cells for the really hard cases. Let the AI handle the easy cases.”

A common theme among responses is the use of human-AI augmenting technologies to improve individual productivity and efficiency by using AI to save time and effort in the workplace. AI strategist I5 specifically brought up insights from a project leveraging AI to help banking advisors do their jobs more effectively by using generative AI to save them time. For example, tools are

implemented to assist with summarizing notes or other long-form content, consolidate sources of data, and use analytics to generate insightful insights about customer relationships.

They outlined a matrix which can be useful when making decisions about which tasks to automate, augment, or retain (Figure 1). By using this matrix, organizations can strategically assess their employee's tasks and retain human involvement in areas where unique human skills and judgement are essential for optimal performance.

Figure 1:
Task Automation-Augmentation Matrix



The matrix categorizes tasks into four quadrants based on their time consumption and level of effort required.

- A. **Low-Time, High-Effort Tasks (Augment):** Tasks that fall into this category require significant effort but can be completed relatively quickly. These tasks are ideal candidates for augmentation through technology to make them easier for employees to perform.
- B. **High-Time, High-Effort Tasks (Retain):** These are tasks that demand both substantial effort and time. They are best suited to be retained by human workers as they involve complex decision-making, creativity, or interpersonal skills that are challenging to replicate for a machine.
- C. **Low-Time, Low-Effort Tasks (Retain):** These tasks are typically quick and easy to accomplish, requiring minimal cognitive or physical exertion. While low-time, low-effort tasks may not be intellectually demanding, they can serve as a source of quick wins and tangible accomplishments for employees.
- D. **High-Time, Low-Effort Tasks (Substitute):** Tasks in this quadrant are characterized by low-effort requirements but consume a significant amount of time to complete.

From this matrix, we see that deskilling via augmentation or substitution is most likely to occur with high-time, low-efforts tasks (such as data entry, basic administrative tasks, or updating a comprehensive spreadsheet) and low-time, high-effort tasks (such as writing a detailed project proposal or developing a complex financial model).

4.3.2. Managerial Control

I3: "I'm not anti-technology – I think technology has an important role to play in improving human society. But so often, under the kind of economic and social system we live in, technologies are developed by companies or by the military in order to serve particular interests. And those interests are not aligned with improving the rights of workers and the power workers have in society."

The idealistic promise of technology was to reduce working hours and free up humans to focus on more creative work while raising living standards and average wealth [8]. However, some interviewees noted that technology is instead being implemented to stifle worker's rights. Specifically, I3 described the rollout of algorithmic management, wherein a growing number of workers are subject to algorithms informing them how to do their jobs rather than forming a close relationship with a manager. They brought up the example of Uber, a company that claimed to be rolling out an innovative and disruptive approach to transportation, but whose real innovation was to attack and decimate the rights that workers had in delivering that service. They mentioned that at the core of Uber's model, beyond the ability to bypass traditional regulation, is an algorithmic system that is shaping and determining how drivers work. Responses from I6 and I7 also indicated that human-AI augmenting technologies are being employed not only to reduce costs, but to increase control over workers.

4.4. Blurred Lines Between Augmentation and Substitution

4.4.1. Augmentation as a First Step Towards Substitution

A common theme among interviews was the lack of clear boundaries between augmentation and substitution, as augmentation can often be the first step towards a full-on substitution strategy. I7 brought up the example of an AI-based tool that was brought in to augment a team of HR managers at JP Morgan Chase in order to identify predictors of a job candidate's potential job performance. After a full year of interaction between the AI tool and the human experts with the goal of mitigating biases in predictions, JP Morgan Chase decided to fully automate away candidate assessment—and lay off a number of managers in the process. The bank's justification for removing humans from this activity was to increase the level of fairness and standardization among candidate assessments while making the process more efficient. In this example, an AI model developed to augment human work ended up being considered sufficiently robust enough to work autonomously without the assistance of a human, thereby showing how the augmentation of certain tasks can eventually lead to the automation of whole jobs altogether.

I5: *“When it comes to augmentation, the top 1% of your workforce is already performing really well because they're doing certain things. What AI does is that it helps the average part of your workforce do the things that your top 1% is already doing, but more easily.”*

Moreover, as AI augments human workers to make them more productive and efficient, fewer workers may be needed as the same level of output can be achieved using fewer inputs. I5 notes that the bank they consult for does not aim to lay people off, but rather to implement technologies and cut the bottom tier of performers due to natural churn in the workplace. So, as a company loses the bottom rung of performers, they will notice that even if they do not replace the workers, they are still able to retain the same level of output as before. However, a company may not hire back for the positions they cut as a result of implementing AI-augmenting technology, indicating the first signs of structural technological unemployment.

4.4.2. Augmentation Eroding the Satisfaction of Human Work

Another complication in human-AI augmentation arises when AI erodes the satisfaction of human work for the sake of improved productivity and efficiency. Specifically, I3 brought up a hypothetical example of a graphic designer who has spent years building up their skills in illustration and typography, to the point of being confident that they are an excellent designer. However, in the precarious world of graphic design, companies augment human work by generating a design using AI and hiring a human to edit it or enhance it to a level that is presentable to the wider public [10]. Thus, companies end up devaluing certain skills by using AI to churn out something that is nearly good enough to use commercially, and later hiring a human to fix the details before they can properly

present it. Under this model, technologies may not primarily be deployed to replace humans, but to change (and often erode) the way human work is done.

5. Recommendations

In terms of recommendations for employers to reduce the negative effects of deskilling in the workplace, two main suggestions were brought up in interviews: 1) to increase the diversity of the teams developing and deploying AI, and 2) to include employees in the decision-making process of human-AI augmentation projects.

Firstly, there is a need to bring in more diverse perspectives when deciding to automate away certain tasks to AI. I3 brought up how the people developing AI tools are engineers and computer scientists who hold a particular view of the world and a particular idea of what skills are important. Conversely, the types of people who are making art and writing fiction full-time are not the types of people who typically sit at the heads of tech companies and help make decisions about what their tools are going to do. Despite recent efforts at increasing diversity in the corporate world, there is still a lack of representation of marginalized groups in the development and deployment of AI [11]. Thus, I8 mentioned the importance of having a team that is diverse “in all senses of the word” to mitigate some of the biases and blind spots that are in place when employing technologies developed by a typical group of software engineers.

Secondly, AI governance researcher I6 mentioned the importance of consulting workers before rolling out an augmentation strategy in the workplace. They recommend holding town hall meetings with one’s workforce and asking employees how AI has changed their day-to-day experiences at work, as well as their fears surrounding AI. An effective strategy I6 has seen is not a top-down or bottom-up approach, but somewhere in the middle where companies can have a conversation with employees while genuinely listening to their concerns. Although I8 highlights that employee consultation is a lengthy process, it allows organizations to gain insights into the employee’s unique needs, challenges, and concerns regarding AI implementation in the workplace to ensure the skills they value are being retained.

6. Conclusion

To conclude, the integration of human-AI augmenting technologies in the workplace has led to an ambiguous landscape of challenges concerning deskilling, technological unemployment, and AI ethics. While it is common for managers to push for an augmentation strategy rather than a substitution strategy, this paper has argued that deskilling can also occur as a result of human-AI augmentation.

Through conducting eight semi-structured interviews with industry professors and AI researchers, two key factors influencing managerial decision-making around automation were uncovered: time/effort (translating to greater productivity, efficiency, and profit) and managerial control. When it comes to augmenting human labour, multiple types of deskilling can occur, from manual to cognitive to emotional. As well, the blurred lines between augmentation and substitution can mean that employers use augmentation as a first step towards a full-on augmentation strategy.

Finally, workers should be wary that the conditions of their work may be remade by a company implementing AI in the workplace to potentially reduce their pay, the rights they have at work, and the power they have in the workplace. While augmentation is peddled as the be-all-end-all solution to societal problems around automation, we must question the notion that implementing workplace technology is an inherent good, and instead think about whether these technologies are being rolled out in a pro-worker way.

Declaration on Generative AI

The author has not employed any Generative AI tools.

References

- [1] Raisch, Sebastian, and Sebastian Krakowski. 2021. "Artificial Intelligence and Management: The Automation–Augmentation Paradox: Academy of Management Review." *Academy of Management Review* 46 (1): 192–210. <https://doi.org/10.5465/amr.2018.0072>.
- [2] Susskind, Daniel. 2021. "Technological Unemployment." In *The Oxford Handbook of AI Governance*, edited by Justin B. Bullock, Yu-Che Chen, Johannes Himmelreich, Valerie M. Hudson, Anton Korinek, Matthew M. Young, and Baobao Zhang, 0. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780197579329.013.42>.
- [3] Tschang, Feichin Ted, and Esteve Almirall. 2021. "Artificial Intelligence as Augmenting Automation: Implications for Employment." *Academy of Management Perspectives* 35 (4): 642–59. <https://doi.org/10.5465/amp.2019.0062>.
- [4] Kim, Tae Wan, and Alan Scheller-Wolf. 2022. "Technological Unemployment, Meaning in Life, Purpose of Business, and the Future of Stakeholders." In *Business and the Ethical Implications of Technology*, edited by Kirsten Martin, Katie Shilton, and Jeffery Smith, 13–31. Cham: Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-18794-0_2.
- [5] Vallor, Shannon. 2015. "Moral Deskillling and Upskilling in a New Machine Age: Reflections on the Ambiguous Future of Character." *Philosophy & Technology* 28 (1): 107–24. <https://doi.org/10.1007/s13347-014-0156-9>.
- [6] Huang, Ming-Hui, Roland Rust, and Vojislav Maksimovic. 2019. "The Feeling Economy: Managing in the Next Generation of Artificial Intelligence (AI)." *California Management Review* 61 (4): 43–65. <https://doi.org/10.1177/0008125619863436>.
- [7] Berdahl, Loleen, and Jason J. Roy. 2021. "Explorations: Conducting Empirical Research in Canadian Political Science." Oxford University Press.
- [8] Hughes, Carl, and Alan Southern. 2019. "The World of Work and the Crisis of Capitalism: Marx and the Fourth Industrial Revolution." *Journal of Classical Sociology* 19 (1): 59–71. <https://doi.org/10.1177/1468795X18810577>
- [9] Bernard Marr. 2023. A Short History of ChatGPT: How We Got To Where We Are Today. *Forbes*. <https://www.forbes.com/sites/bernardmarr/2023/05/19/a-short-history-of-chatgpt-how-we-got-to-where-we-are-today/?sh=286b1544674f>
- [10] Woodruff, Allison, Renee Shelby, Patrick Gage Kelley, Steven Rousso-Schindler, Jamila Smith-Loud, and Lauren Wilcox. 2024. "How Knowledge Workers Think Generative AI Will (Not) Transform Their Industries." In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–26. Honolulu HI USA: ACM. <https://doi.org/10.1145/3613904.3642700>.
- [11] Collett, Clementine, and Sarah Dillon. 2019. "AI and Gender: Four Proposals for Future Research." *Apollo - University of Cambridge Repository*. <https://doi.org/10.17863/CAM.41459>.
- [12] Keynes, John Maynard. 1930. "Economic Possibilities for Our Grandchildren." In *Essays in Persuasion*, edited by John Maynard Keynes, 321–32. London: Palgrave Macmillan UK. https://doi.org/10.1007/978-1-349-59072-8_25.