PARTNERSHIP ON AI

WHITE PAPER

AI and Job Quality

Insights from Workers

Stephanie A. Bell

SEPTEMBER 13, 2022



Contents

Executive Summary **4**

Introduction **6**

The need for workers' perspectives on workplace AI 6

The contributions of this report 8

Our approach 9

- Key research questions 9
- Research methods 9
- Site selection 9
- Who we learned from 11
- Participant recruitment 12

Major Themes and Findings **13**

Theme 1: Executive and managerial decisions shape AI's impacts on workers, for better and worse 13

Theme 2: Workers appreciate how some uses of AI have positively changed their jobs 16

Theme 3: Workplace AI harms repeat, continue, or intensify known possible harms from earlier technologies 18

Theme 4: Current implementations of AI in work are reducing workers' opportunities for autonomy, judgment, empathy, and creativity 20

Theme 5: Empowering workers early in AI development and implementation increases opportunities to implement AI that benefits workers as well as their employers 22

Opportunities for Impact **24**

Stakeholder 1: AI-implementing companies 25

Stakeholder 2: AI-creating companies 26

Stakeholder 3: Workers, unions, and worker organizers 28

Stakeholder 4: Policymakers 29

Stakeholder 5: Investors 31

Conclusion **32**

Acknowledgments **33**

The AI and Shared Prosperity Initiative Steering Committee **34**

Appendices **35**

Appendix 1: Detailed site and technology descriptions 36

Appendix 2: Research methodology 40

References **44**

LIST OF FIGURES

Figure 1 Research & insight creation approach 9

Figure 2 Stages of employee lifecycle/worker journey where AI tools may be implemented 10

Figure 3 Participant demographics & characteristics: Overall 36

Figure 4 AI technology types included in study by worker journey stage and site 36

Figure 5 Participant demographics & characteristics: India 37

Figure 6 Participant demographics & characteristics: Sub-Saharan Africa 38

Figure 7 Participant demographics & characteristics: United States 39

Executive Summary

Across industries and around the world, AI is changing work. In the coming years, this rapidly advancing technology has the potential to fundamentally reshape humanity's relationship with labor. As highlighted by previous Partnership on AI (PAI) research, however, the development and deployment of workplace AI often lacks input from an essential group of experts: the people who directly interact with these systems in their jobs.

Bringing the perspectives of workers into this conversation is both a moral and pragmatic imperative. Despite the direct impact of workplace AI on them, workers rarely have direct influence in AI's creation or decisions about its implementation. This neglect raises clear concerns about unforeseen or overlooked negative impacts on workers. It also undermines the optimal use of AI from a corporate perspective.

This PAI report, based on an international study of on-the-job experiences with AI, seeks to address this gap. Through journals and interviews, workers in India, sub-Saharan Africa, and the United States shared their stories about workplace AI. From their reflections, PAI identified five common themes:

1. Executive and managerial decisions shape AI's impacts on workers, for better and worse. This starts with decisions about business models and operating models, continues through technology acquisitions and implementations, and finally manifests in direct impacts to workers.
2. Workers have a genuine appreciation for some aspects of AI in their work and how it helps them in their jobs. Their spotlights here point the way to more mutually beneficial approaches to workplace AI.
3. Workplace AI's harms are not new or novel — they are repetitions or extensions of harms from earlier technologies and, as such, should be possible to anticipate, mitigate, and eliminate.
4. Current implementations of AI often serve to reduce workers' ability to exercise their human skills and talents. Skills like judgment, empathy, and creativity are heavily constrained in these implementations. To the extent that the future of AI is intended to increase humans' ability to use these talents, the present of AI is sending many workers in the opposite direction.
5. Empowering workers early in AI development and implementation increases the opportunities to attain the aforementioned benefits and avoid the harms. Workers' deep experience in their own roles means they should be treated as subject-matter experts throughout the design and implementation process.

In addition, PAI drew from these themes to offer opportunities for impact for the major stakeholders in this space:

1. AI-implementing companies, who can commit to AI deployments that do not decrease employee job quality.
2. AI-creating companies, who can center worker well-being and participation in their values, practices, and product designs.
3. Workers, unions, and worker organizers, who can work to influence influence and participate in decisions about technology purchases and implementations.
4. Policymakers, who can shape the environments in which AI products are developed, sold, and implemented.
5. Investors, who can account for the downside risks posed by practices harmful to workers and the potential value created by worker-friendly technologies.

The actions of each of these groups have the potential to both increase the prosperity enabled by AI technologies and share it more broadly. Together, we can steer AI in a direction that ensures it will benefit workers and society as a whole.

Together, we can steer AI in a direction that ensures it will benefit workers and society as a whole.

Introduction

The need for workers' perspectives on workplace AI

In the past decade, global investment in artificial intelligence development has soared. Private investment in AI went from under $5 billion globally in 2013 to over $90 billion in 2021, more than doubling between 2020 and 2021 alone.[1] The implementation of AI products has similarly grown: In 2021, 56% of respondents to a McKinsey survey said their organizations used AI in at least one business function[2] compared to 20% of respondents in 2017 who reported using AI at scale or in a core part of their business.[3] The positive and negative effects of this are already being felt by both formal workers (millions of whom are interacting with AI products or will soon see them incorporated into their jobs) and informal workers (who are encountering transformed market conditions due to the use of AI by businesses). For both groups of workers, the positive and negative impacts of these technologies are unevenly distributed, often following other existing axes of inequality, such as geography, race, and gender. Yet workers' needs, well-being, and expertise are under-considered in AI research, development, and implementation.

In an earlier publication, "Redesigning AI for Shared for Prosperity: An Agenda,"[4] PAI highlighted the need to better understand AI's impacts on job quality, including by engaging the workers who experience these impacts firsthand. Workers who directly interact with AI understand these systems' benefits and harms in depth. In the best of circumstances, they experience the ways these technologies can make their work more efficient, error-free, and pleasurable or less grueling, tiring, or dangerous. Too frequently, workers also experience the downsides. These systems can restrict workers' autonomy, invade their privacy, undercut their judgment and empathy, and push them to the point of exhaustion or injury. Companies that allocate managerial tasks to AI systems can subject workers to binding decisions that are capricious or cruel.

At a societal level, the increasing adoption of AI systems is poised to accelerate existing problems arising from economic inequality.[5] AI research and product development is taking place in a highly concentrated group of countries and companies. Private AI investment in the United States in 2021 totaled $52.9 billion, over three times the investment by the next highest country, China at $17.2 billion — which in turn exceeded investment by the next nine countries combined.[6] The impacts of workplace AI use, however, will be felt around the world. As some companies attempt to automate work they had previously outsourced, others will adopt AI systems created in and for entirely different geographies.[7] Both between and within countries, AI's current trajectory threatens to widen the gaps between the haves and have-nots.

Moreover, workers are uniquely positioned to understand how to avoid these harms and contribute ideas to improve their employers' bottom lines. Using AI to increase job quality (or at least not decrease it) would enable employers to reap the benefits of a more engaged and satisfied workforce. Higher job quality and employee satisfaction increases productivity of existing workers, reduces turnover (retaining experience and expertise), and fosters the ability to recruit higher-caliber talent in competitive labor markets.[8] Decades of research on innovation in domains as diverse as manufacturing,[9] [10] hospitality,[11] [12] and government service provision[13] has underscored the unique insights and innovative potential of frontline workers and other individual contributors. Workers are afforded intimate knowledge of crucial aspects of their work that managers and leaders only see from a distance. They are experts in things like the nuances of how to create the conditions for customer satisfaction or the levels of care that need to be taken in moving objects of different fragility through a warehouse. This deep knowledge of the ins and outs of completing core tasks makes workers an underutilized source of expertise on issues and problems where AI could be a powerful tool or assistant.

Workers are uniquely positioned to understand how to avoid AI's harms and contribute ideas to improve their employers' bottom lines.

WHY WON'T THE MARKET ADDRESS HARMS BY INCREASING WAGES?

Strict rationalist economic theory would predict that workers will receive sufficient wages to compensate for technologically driven harms.[14] However, employers and workers alike lack the perfect information required for this effect. Additionally, this theory presumes robust competition for labor, and workers who possess a genuine ability to choose between different employment options.

In many labor markets, employment options are relatively concentrated, enabling companies to treat workers worse than they would in more competitive environments.[15] [16] [17] Steps to increase information and awareness can reduce the likelihood that workers unwittingly accept poor working conditions without a sufficient compensating wage. Regulation and increased unionization can reduce the negative effects of concentrated labor markets. However, the insufficiency (as well as improbability) of these solutions point to a need for direct attention to AI's effects on job quality.

Finally, pursuing collaborative workplace AI that draws on the unique strengths of humans and technology enables businesses to expand the production frontier. Many current integrations of AI into human workflows are designed around the limited capabilities of the AI systems. This, in turn, circumscribes the range of talents and skills of the people who work with them. Starting from the opposite premise — that AI should be integrated into workplaces in a way that enables human skills and talents to flourish — is undeniably harder. The reward for the achievement, however, is far greater for both workers and their employers.

The contributions of this report

Past research and discussions on AI's impacts on workers have frequently taken one of three forms, with the first two especially common in popular and business discussions:

1. Predictions about AI's impacts on job availability (i.e., how many jobs AI will eliminate and which ones).
2. Aspirational discussions of how AI will improve work for humans by automating "dull, dirty, and dangerous" work.
3. Targeted research by academics and civil society groups on the negative impacts of AI focused on specific technologies or groups of workers.

In this last category, groundbreaking research has illuminated the harms of specific AI technologies and use cases,[18] including monitoring and surveillance,[19] [20] algorithmic decision-making,[21] [22] shift-scheduling,[23] and platform work software and applications.[24] [25] Researchers have also explored particular types of impacts on workers, including worker health and safety,[26] data collection and privacy,[27] [28] [29] and reproductions of carceral power.[30] Previously, PAI itself conducted a landscape review of AI's demonstrated and potential impacts on worker well-being.[31]

This report builds on this foundational work by bringing in the perspectives and experiences of frontline workers at the frontier of workplace AI implementation around the world. It shares their stories of how their jobs have been transformed by AI (for better and for worse) and highlights their oft-neglected expertise on challenges and opportunities in their work where they welcome AI assistance. It also synthesizes this primary research with the existing literature to offer implications and opportunities for key stakeholders on how they can take action to ensure the category of technological products commonly referred to as AI improves — not worsens — the experience of workers. Finally, it offers areas in need of further exploration in future research or implementation case studies.

Through their comments and stories, workers surfaced five key themes about their experiences of AI in the workplace. These five themes point the way toward a better future for workplace AI, one that maintains or increases companies' profitability and revenue while also maintaining or increasing job quality. Getting there will require many decision-makers and stakeholders to do things differently than they have in the past. In some instances, the needed changes are substantial and complex. At the end of this report, we offer initial recommendations for all of the major stakeholders in this space: AI-using companies, AI-creating companies, workers and the organizations such as unions that represent them, policymakers, and investors.

Our approach

Key research questions

We set out to understand workers' experience of AI in their jobs and possible opportunities to foster worker participation and voice in the processes of AI creation and deployment. Our key research questions were:

- How and why are AI and automation technologies changing workers' tasks, coaching, and evaluation? What are workers' reflections on the positives and negatives of those changes?
- How is workplace AI affecting different aspects of worker well-being (including purpose, meaning, and physical, emotional, and intellectual well-being)[32] as articulated and valued by workers themselves?
- In what ways are workers currently participating in the creation and implementation of AI used in their workplaces? How much direct influence or decision-making power do they see themselves as having in these processes? Would they be interested in more ways to participate? Why or why not?

Research methods*

* Additional details on research methods can be found in Appendix 2.

The findings of this report are grounded in a review of the existing research and two types of primary qualitative research we conducted with workers: diary studies and interviews. All primary research was conducted virtually due to the pandemic.

FIG. 1 Research & insight creation approach



1. Literature review
Review existing academic and civil society research, identify impact areas with low coverage & opportunities to add actionable detail

2. Diary studies
Asynchronously solicit near-real-time, less-mediated storytelling from participants on how AI has transformed their jobs

3. Interviews
Conduct semi-structured interviews to expand upon and test emerging themes, explore outlier responses and views

4. Analysis
Deductive technology use clustering, inductive job quality theme analysis

Site selection*

* Additional details on sites, included technologies, and participant demographics can be found in Appendix 1.

AI's transformation of work is global in scope. Its impacts within and across countries often follow deeply grooved paths of inequality created by past economic and political systems.[33 34 35] We conducted the primary research with three groups of workers, focusing on people holding individual contributor (rather than managerial) roles:

- Customer service agents in India
- Data annotators in sub-Saharan Africa
- Warehouse workers (e.g., pickers, packers, loaders) in the United States

We sought a range of geographies, industries, occupations, and skills. This multisite approach allowed us to explore diverse experiences on issues including:

- Whether there is something inherent to artificial intelligence as a technology (irrespective of geography, industry, and worker skills) in how it transforms work
- The economic opportunities and vulnerabilities associated with varying wage and income levels (between different occupations and geographies)
- Worker attitudes toward jobs and labor
- Labor market structures and near-term susceptibility to technological disruption by AI
- Managerial decision-making cultures
- Government interest in fostering local AI ecosystems
- Government interest and capacity to regulate AI's impacts on workers[36] [37]

FIG. 2 Stages of employee lifecycle/worker journey where AI tools may be implemented



While workers may be affected by AI technologies across the "employee lifecycle,"[38] this research focuses on use cases where workers could directly observe AI in their workplace and thus share their experiences of it.* In line with this approach, this report also discusses different technologies as experienced by the participating workers, in what might be thought of as the "worker's journey" in a given job. As an illustration, a technology used by an employer to provide guidance to workers on their tasks is discussed as a worker's coaching tool, not as a manager's automation or decision-support tool.

Workers largely encountered the AI technologies described in this study in three stages of their jobs: in their fulfillment of their roles, to coach them on their work, or to evaluate their performance.

* Technologies used in other stages, such as AI recruitment or assessment tools for job candidates, pose their own risks to workers; for an overview of a number of these areas, see "Platforms at Work: Automated Hiring Platforms and Other New Intermediaries in the Organization of Work."[39]

An area that has seen substantial research is the ways these tools reinforce biases related to race, gender, age, class, and disability status. Compared to many aspects of workplace AI, this area has received greater attention from more traditional AI ethics areas, such as fairness, transparency, and accountability, as well as from major employers. For an overview of potential AI bias issues in hiring, see "All the Ways Hiring Algorithms Can Introduce Bias."[40]

For an effort undertaken by employers to address issues related to algorithmic bias in workforce decisions, see "Algorithmic Bias Safeguards for Workforce."[41]

For a tool to assess potential issues with automated hiring software, see "A Silicon Valley Love Triangle: Hiring Algorithms, Pseudo-Science, and the Quest for Auditability."[42]

WHY WORRY ABOUT THE QUALITY OF JOBS THAT ARE TARGETED FOR AUTOMATION?

Substantial research and capital is being dedicated to automate the jobs covered in this research. Automation aspirations, however, are not automation realities.

Since the invention of automated retrieval systems in the 1950s, technology has enabled warehouse operators to process goods at ever-increasing speed. Yet more than 1.8 million people in the US work in the warehousing and storage industry today, more than 2.5 times the number a decade ago.[43] Efficiency gains have lowered costs and increased consumer demand, but cost-effective, fully automated warehouses remain elusive for most goods. A similar narrative holds for call centers, with speculation about automating customer service since the creation of the ELIZA chatbot in the 1960s. Despite the conversion of recent natural language processing research into cutting-edge call center software, global demand for call center workers remains strong and the general public remains skeptical of fully automated customer service.[44] Automation efforts often transform the tasks of jobs. But the demand for workers in both occupations continues, with no reliable endpoint in the near or medium term.

Data annotation is a comparatively new "job gained" from the rise of AI.[45] It, too, is an automation target. As with the other jobs, automation is changing the content of the role, but demand for workers continues to grow, as increasingly sophisticated AI-modeling techniques demand new forms of data.[46]

Furthermore, AI technologies similar to those in this report are proliferating throughout other jobs of all kinds. Without deliberate intervention, the decisions, incentives, and technologies underpinning the impacts discussed in this report will likely reproduce similar impacts for workers in other industries and locales. The occupations discussed here may not survive indefinitely — few jobs do. The aim of this report is not to protect specific jobs, but rather to address the interests of workers broadly (especially economically vulnerable workers) in the face of rapid technological change.

Who we learned from

The occupations and locations selected feature workers representing a diversity of profiles. In India, offshore call center work is a stable, middle-class job often performed by college-educated workers fluent in a second language (in this case, English). Shifts are scheduled according to the needs of the outsourcing country, so workers often find themselves working night shifts, very early, or very late to match standard working hours in high-income English-speaking countries around the world. These hours are hard to navigate alongside family life; it is estimated that over 90% of people in these roles are under the age of 35.[47]

The data annotation workers in the sub-Saharan country where this research was conducted are similarly youthful. This work is often positioned as an entry-level job for those interested in the continent's growing information technology industry. As educational requirements are less strict than those for offshore customer service workers in India, educational backgrounds are more varied. Most workers have at least a secondary school degree and many have gone on to take classes in or complete post-secondary or bachelor's degrees.

Unlike in the other two sites, the demographic profile of warehouse workers in the United States is highly heterogeneous. The purpose of the work means worksites are distributed throughout the country rather than concentrated in a handful of cities or a region. Substantial skill or education requirements are uncommon for entry-level jobs in the industry. Given the US's history of education inequality and the lack of access to quality education for many people of color, people of color are overrepresented in warehouse work in the US, making up nearly 60% of the industry workforce.[48]

Participant recruitment*

* Additional details on participant recruitment can be found in Appendix 2.

Participants for this research were recruited in two ways:

INDIA AND US

We worked with a professional recruiter to identify interested and qualified candidates for the study. Participant groups for each site were then selected from these pools to create samples representative of the demographics of workers in those occupations in those locations.

SUB-SAHARAN AFRICA

Participants from the sub-Saharan Africa site work for a company that had developed a set of machine learning applications to assist their employees. This company was interested in better understanding its employees' perspectives on the new software and the ways it has changed their work. The company facilitated introductions to a group of employees with experience using the software. The group could opt into the research. Participation was entirely voluntary. A strict firewall was implemented from the outset of the research to protect participant confidentiality and ensure they felt comfortable speaking freely about their experiences.

HOW IS AI AFFECTING INFORMAL WORK?

Though the primary research in this report is focused on formal sector workers, 60% of the world's workers participate in the informal sector.[49] While few, if any, work directly with AI systems, AI is still transforming their work by changing informal labor market conditions.

Consider agriculture, where over 90% of the workforce is informal.[50] Globally, informal workers are two times more likely than formal sector workers to be members of the working poor and agricultural workers are more likely than other informal workers to be poor.[51] Many are sharecroppers and contract farmers who make deals with formally incorporated companies to grow specific quantities of specific crops over a given period of time. Historically, negotiations would take into account an informal worker's accumulated knowledge of local soil and weather conditions, performance of past crops, prior market prices, and other factors. Informal farmers possess the type of experiential knowledge passed through communities and generations, which can be formidably accurate.[52] The companies, on the other hand, possess the type of technocratic knowledge built through the collection and increasingly sophisticated analysis of data.

With the introduction of AI, informal farmers' ability to negotiate critical provisions has radically decreased. Many farmers now face take-it-or-leave-it offers to produce crops they've never seen grown and which may require a year or more of invested cultivation before producing sellable yields. The unprecedented nature of the offers means farmers lack experiential knowledge to base their decisions on, and contracting companies are not sharing the details underlying their proposals, creating sizable information asymmetries.

The experience of the Self Employed Women's Association (SEWA) in India working with women farmers in the informal sector has revealed a lack of inclusive, quality data. The algorithms used by companies rely on data collected by researchers and economists. Informal sector workers, and in particular women, have not been included in the design of data collection tools or the data collection itself. The exclusion of their perspectives and knowledge raises questions about the usability, authenticity, and relevance of this data.

The contracts created with this data are increasingly non-inclusive, and transfer risks to poor smallholder farmers, pushing them deeper into a vicious circle of poor data representation, poor contracts, high risk, increased poverty, and ever-growing debt. There needs to be a substantial focus on including small and marginal women farmers in the data collection processes, resulting in transparent and inclusive data captured firsthand from informal sector workers.[53]

Major Themes and Findings

The workers participating in this research shared stories, experiences, and observations of their time interacting with AI in their workplaces. The findings of this report are drawn from their insights. While common themes manifested differently according to setting, they appeared across all of the research sites and reflect what we heard from a substantial portion of participating workers. The ways these themes might present themselves (including in settings beyond those we researched) depend on a number of factors, including regulatory protections, companies' managerial priorities, and workers' relative influence in their workplaces (through unions, worker organizations, or individual leverage due to local labor market conditions). Workers also experience these impacts unevenly as individuals. Personal demographic characteristics — such as their race, age, gender, immigration status, disability status, and formal education level — may lead them to be more marginalized or vulnerable.

THEME 1 Executive and managerial decisions shape AI's impacts on workers, for better and worse

Workplace AI is deployed by particular executives and managers in specific contexts and specific ways. Leaders and managers determine whether to use workplace AI technologies, which workplace AI technologies to use, what goals they are intended to accomplish and how they are to be used. These decisions are driven by a combination of business models, company culture, industry trends, and the availability of relevant AI products. These factors also shape each other. The initial choices made by companies that produce technological impacts on their workers are, at first glance, not technology decisions at all: they're foundational choices about the operating model and personnel strategy of their business.

How hierarchical is the business model? How much discretion are employees given to use their own judgment in executing their work (as opposed to following a strict set of rules and procedures)? Are employees encouraged to stay and develop expertise and experience that they bring to their roles or intentionally churned to keep costs low? Are jobs designed so that they can be performed with very little training (rendering workers intentionally interchangeable) or do they reward experience? How aggressively are performance targets pushed and punished?

These foundational decisions in turn structure subsequent decisions about what technologies could be useful in meeting business goals, as well as how they ought to be used. Upstream decisions on questions like these likely have a significantly higher influence over how AI affects workers than any choices made by their immediate managers. For instance, a company that designs a "high road" model and strategy to offer its workers high degrees of autonomy (a job attribute highly correlated with high job quality and employee satisfaction[54])

would likely see more value in non-binding AI decision-support tools. On the other hand, a company that designs a "low road" model, with its roles to be performed with very little training or autonomy and very short average tenure (highly correlated with low job quality and employee satisfaction[55]) would see more use in technologies that closely monitor work to ensure it is being performed correctly or claim to remove the need for human judgment. Each of these decisions has an impact on workers, shaping what tasks they are expected to accomplish and how they are expected to do them. All of these decisions impact workers beyond technology, potentially much more than any technologies used — but they also shape workplace AI's impacts.

The initial choices made by companies that produce technological impacts on their workers are, at first glance, not technology decisions at all: they're foundational choices about the operating model and personnel strategy of their business.

As an example, the customer service agents we spoke to in India use AI software marketed to customer support companies and teams as real-time coaching, performance assessment, and task augmentation for their agents. One function of the software is to monitor their calls and text chats for keywords and phrases to diagnose possible customer issues and suggest resolutions, which are offered to agents in real-time pop-ups and menus. Another function is to monitor tone of voice, volume, and keywords to assess emotion and offer real-time pop-ups and alerts to agents on how they could better manage the emotional side of their interactions with customers (for instance, warnings that conversations are sounding emotionally charged, or guidance to speak more quietly, or slow down their speaking speed).

In the agents' use of this software, two clear examples of this theme emerged. First, for some agents, it was clear from their employers' guidance that they should take AI alerts and prompts they received during calls or text chat sessions as suggestions, rather than commands or requirements. This group of workers was expected to exercise autonomy and judgment in meeting customer needs, using the AI feedback as one of many inputs in their call or chat handling. Agents at different companies were expected to closely follow the feedback from the AI and not disregard its recommendations except, perhaps, in extreme circumstances. Both groups recounted instances where they judged the AI to be incorrect in its recommendations, but the group empowered to deploy their judgment on calls or in chats felt more autonomy and control over the quality of the service they provided. Some employers treated AI feedback and call assessment (including predictions of customer feedback scores) as purely a coaching tool. Others used it as a direct input to performance evaluations. In a coaching setting, workers were able to put the feedback in context for themselves, adopting suggestions where they made sense. In a performance evaluation setting, the context was often flattened or missing, adding an element of arbitrariness where managers likely intended to add rigor.[56]

In sub-Saharan Africa, the data annotators were tasked with annotating images or videos for data sets to be used in developing machine learning (ML) models. Prior to the introduction

of machine learning software to automate part of this work, the workers carefully outlined each contour of an object in an image. For videos, this could require them to meticulously shift the position and edit the contours of the outline for dozens or even hundreds of frames where the object had only slightly changed position from one frame to the next. The company recently introduced ML task automation software to assist the workers in the fulfillment of their roles. For certain objects in an image, workers could identify the outermost corners of the object and the software filled in the rest. For videos, the software could take the initial object outline delineated by the worker and predict the outlines of that object in many future frames.

The workers who participated in this research were tasked with testing and providing feedback on the company's new ML software (in addition to being responsible for actual annotation work). Unlike other workers responsible for specific client deliverables and deadlines, they were not given strict quality or completion performance targets for the annotation side of their role. They still, however, had the opportunity to earn bonuses for the speed and accuracy of their work. This incentive structure for their work gave them the needed time and freedom to focus and reflect on improvements to the ML tools that could deliver value for the company without forcing them to miss out on the opportunities for additional compensation offered to their colleagues.

Previous research has offered other demonstrations of how managerial decision-making shapes AI technology's impacts on workers.[57] This includes the use of big data analytics as invasive and harmful "bossware,"[58] the cruelty that can result from algorithmic decisions with no human recourse,[59] [60] the negative health impacts of overly aggressive performance targets set using AI,[61] [62] the lack of worker protections afforded to workers misclassified by their employers as independent contractors on AI-driven platforms,[63] [64] and the negative health impacts, life disruptions, anxiety, and job insecurity arising from last minute shift-scheduling enabled by AI software.[65]

Better choices are available for managers and leaders implementing AI in their workplaces.

These negative impacts on workers should not be seen as inevitabilities of the unstoppable march of technological progress, but rather as the outcomes of a series of decisions. These decisions are made first by companies who create business and operating models revolving around low-quality jobs, then by product developers and designers who build AI technologies that are either explicitly designed for these uses or possible to misuse in harmful ways, and subsequently by leaders and managers who choose these particular implementations. The beneficial examples above, where AI software was used to assist workers while they maintained their autonomy and retained decision-making authority, demonstrate that better choices are available for managers and leaders implementing AI in their workplaces.

THEME 2

Workers appreciate how some uses of AI have positively changed their jobs

While there are clear harms arising from some workplace AI uses and decisions, the role of workplace AI in job quality is not wholly negative. Across our research sites, workers offered reflections on what they appreciated about specific uses of AI or attributes of AI products that they use. In the India site, in addition to the appreciation for additional information to support their decisions, and real-time coaching that was not evaluative or punitive, workers highlighted time-saving as well as benefits to their physical well-being from AI software that logged caller details and auto-prompted solution menus. The call center workers also reported that the software's automated data entry reduced eye strain and repetitive stress injuries to their wrists and hands compared to the constant keyboard, mouse, and screen work needed when entering this information themselves.

In the sub-Saharan Africa site, a strong majority of the data annotators preferred working with the ML tools compared to when they did their work more manually. They lauded the speed with which the ML prediction software enabled them to complete annotation tasks, and the reduction in sometimes tedious or boring repetitive work. (For instance, working their way through each frame in a video from start to finish.) Some workers also mentioned that the tool helped them feel less tired throughout the day or at the end of their shifts.

However, they noted the software also sometimes had accuracy problems. In these cases, many workers would have preferred to manually complete those tasks themselves from start to finish. When the software was inaccurate in its outputs, it posed several problems to the workers. First, it forced them to use their time inefficiently — not only did they have to spend the time waiting for the algorithm to complete its (incorrect) annotation, they also had to spend additional time revising the output from the software. Second, the process of trying to find each error and then correct it felt unnatural and painstaking compared to when they felt mentally prepared to just do the tasks themselves. Finally, they felt a sense of frustration familiar to anyone required to work with a malfunctioning technology: the software was failing to meet their expectations and leaving them to sort out the problems it created. In interviews, the data annotators explained that part of this performance gap could be attributed to portrayals of the technology when it was introduced. Because it was "machine learning" or "artificial intelligence," they expected it to be more accurate than their own work, not less. Still, even workers who expressed these issues praised the benefits listed above when the technology was working properly.

While there is a broader, ongoing discussion of puffery in the AI industry, less coverage has been afforded to the effects of similar dynamics in workplaces.[66] [67] Inflated portrayals of workplace AI's capabilities may do more harm than good. Setting high expectations (however inadvertently) and then failing to meet them was a source of frustration and stress expressed by the call center workers regarding the call-coaching and evaluation software

as well. In the context of AI's benefits to workers, setting realistic expectations and then meeting or exceeding them may substantially reduce friction in AI use.

Among the warehouse workers in the US, many singled out AI technologies that reduced possible errors, such as placing items in the wrong locations or using the wrong tape or labels on packages. A number of participants said they felt an increased degree of pride in their job due to their accuracy in their work. Some research participants additionally valued how warehouse robots reduced some physical demands of the job. In the case of robots that bring items to workers, this could be a radical reduction in steps walked by workers who previously would have walked 10 or 20 miles a day to get these items themselves. The assistance of robotic arms could reduce muscle strains and pulls. Positive reactions to these physical effects were mixed, however, with some participants noting that they missed the exercise they got in the old way of working and others raising concerns about increases in injuries from repetitive movements prompted by the robot-assisted workflow.[68] [69] [70]

The workers who spoke with us were not anti-technology or anti-AI. They welcomed technological assistance, provided they could maintain or improve their job quality while using them.

Some of these benefits of workplace AI commonly cited by workers — like increases in speed, accuracy, efficiency, and productivity — were clearly intended by the creators and implementers of the technology. Others, such as the sense of pride in a job well done, could be seen as indirect effects of those benefits intentionally sought by the AI creators and implementers. Still others, such as the ergonomic advantages of automated call-logging, were meaningful improvements to worker well-being that likely did not play a decisive role in the creation of the software or the company's decision to purchase it, but accrued to the worker nonetheless.

Both the intended and unintended positive consequences of workplace AI cited by workers point towards possible paths for developing and implementing workplace AI technologies that benefit workers as well as their employers. The workers who spoke with us and shared their stories and experiences were not anti-technology or anti-AI. Their own descriptions of what counts as a good work day and their personal definitions of what it means to do good work share a number of values and goals with their employers, including swift and accurate completion of their tasks. The participating workers welcomed technological assistance in achieving these objectives, provided they could maintain or improve their job quality while using it. The positive experiences of AI and perspectives shared by workers should give businesses confidence that benefits to workers and benefits to employers are not zero-sum. Workplace AI integrations can deliver value to both groups. Respectful, considered AI implementations that maintain worker dignity and autonomy can be embraced by workers.

THEME 3

Workplace AI harms repeat, continue, or intensify known possible harms from earlier technologies

While AI may be relatively new to most workplaces, the impacts workers see from its use are not. Many negative impacts from workplace AI are versions of impacts seen from non-AI technologies. For instance, employers may make job and task design decisions encouraging repetitive motions that could lead to injuries (as reported by some participating warehouse workers) in order to integrate AI task automation into workflows. This also occurs in other, non-AI assisted industrial settings where workers are assigned a small set of tasks to perform repeatedly. AI systems can deliver negative feedback to workers without helpful suggestions for improvement: an issue noted by some participating customer service agents and also an unfortunate practice of some human managers since the creation of managerial and supervisory roles. Additionally, some companies deployed intensive monitoring of their workers well before big data and AI made it possible for managers to analyze that data in increasingly invasive and stressful ways.

A US warehouse worker offered a representative explanation of how a performance evaluation AI system layered into her job — a monitoring software used by her company to provide real-time performance feedback — negatively affected her emotional well-being. From when she clocks in until she clocks out, she is constantly monitored by software. The software tracks when she is completing a task (for instance, following instructions she has been given about how to process an item in the warehouse). It tracks how long it takes her to complete that task, tracks when she is between tasks, tracks when she goes to get water or use the bathroom. And it tells her whether she is staying on pace or falling behind the goals her company managers use that same data to set. The expectation is that she is constantly on pace. If she falls behind for any reason, it triggers stress that stays with her until she is ahead of the targets again. The stress isn't from personal perfectionism: firing is a common consequence for workers who fall behind targets at her employer, regardless of whether they might have understandable reasons for a slower pace (for instance, health conditions that might require more frequent breaks).

The pressure generated by the way her company management uses this software leads her and her colleagues to cut corners to speed up their work. When they're trying to stay on pace, she and some of the other warehouse workers pointed out that they would sacrifice safe or proper movements or lifting techniques in favor of speed. The consequences their employers set for being too slow made the choice clear for them: they focused on not getting fired over making sure they stayed safe. While some technologies in AI-assisted warehouses can reduce physical burdens on workers, such as robotic item movers which reduce the distances workers walk in a shift, employers' decisions to use AI technologies to accelerate the pace of work can create higher worker injury rates.[71][72]

On top of the emotional and physical well-being issues that workplace AI can cause, the way managers and executives choose to integrate AI into the overall workflow may lead to lowered intellectual well-being on the job. Workers at each site largely agreed that the AI systems used in their jobs lowered the level of intellectual challenge when compared to what it looked like to do their work without AI. Workers in US warehouses with higher degrees of AI implementation often had less variety in their tasks and more technological guardrails to assist them in performing them correctly. The customer service agents in India spent less time and energy diagnosing the reasons a customer called or identifying possible solutions for their issues. In sub-Saharan Africa, the data annotators no longer completed intricate tasks requiring a careful, discerning eye from start to finish, but instead largely spent their time creating broad outlines around objects in images, letting algorithms do the rest. While many welcomed the extra ease, many others indicated that they preferred a higher degree of challenge.

Workers at each site largely agreed that the AI systems used in their jobs lowered the level of intellectual challenge. While many welcomed the extra ease, many others indicated that they preferred a higher degree of challenge.

Each of the examples offered above can be seen as a continuation of trends from other workplace technologies. However, existing laws and regulations do not appropriately address these harms. The status quo enforcement of basic health and safety protections for workers around the world is inadequate to prevent them from being harmed by their jobs: the introduction of AI software and systems that can ratchet up work intensity only increases the urgency of shoring up these laws and their enforcement.[73] In addition to the emotional and mental health impacts described above, AI monitoring and surveillance technologies undermine workers' sense of privacy, dignity, and autonomy.[74] Yet mental health safeguards are often less regulated or enforced and a policy vacuum exists in many geographies regarding privacy and data protections at work.

The familiarity and continuity of harms from workplace AI should make them easier to anticipate, and thus to prevent or mitigate through responsible design and use. But until consideration of these impacts is foregrounded by AI developers and the executives and managers who purchase and implement workplace AI, or sufficient protections and enforcement are enacted by governments, workers will continue to suffer harms that could have been anticipated and prevented.

THEME 4

Current implementations of AI in work are reducing workers' opportunities for autonomy, judgment, empathy, and creativity

One optimistic line of thought on AI's transformational effects on jobs suggests AI will support humanity to take on more creative, empathetic, or intellectually advanced work[75] [76] — an admirable goal of creating jobs with more "human" tasks than the reproducible, mathematically definable work of algorithms and robots. As reported by workers collaborating closely with AI, however, the current reality on the ground points to jobs moving in the opposite direction.[77]

Take the data annotators working with an ML technology that automates some of the annotation work they would have previously done manually. Their responsibilities shifted away from a creative, generative role that some of them described as like a craft or art. Previously, they carefully drew the outlines of relevant objects and derived satisfaction from their precise handiwork. With the addition of ML software to their workflow, they now, in their words, spend less time creating and more time "fixing" or "cleaning" the AI's output by identifying and editing images that the algorithms annotated incorrectly.

Some of the call center workers used technologies designed to address two of these supposedly more human skills: empathy and problem-solving. For empathy, call monitoring software assessed whether calls were getting too emotionally charged by measuring agents' volume, speed, and word choice. For problem-solving, a software designed to assist agents detected keywords and phrases in order to pull up solution lists and suggest possible issues the customer or client might be having. In each case, the software was not consistently accurate or helpful (according to the judgment of the experienced call center workers) but workers often had to contend with performance assessments tied to complying with this software by making their displays of empathy more templatized and, ultimately, less human.

In automated warehouses, workers who had been around prior to the introduction of new AI and robotics systems, or who switched from warehouses with less automation to cutting-edge robotics locations, found that the variety of their tasks shrank over time, with AI, robotics, and other automated systems picking up tasks that they previously completed or coordinated with other workers to complete. Multiple workers mentioned feeling like they themselves were also robots in the more automated warehouses. Along with workers' increasingly parochial view of the work being done throughout the warehouse came a decrease in their positioning and ability to identify and suggest improvements at a systemic level. Their universe of problem-solving potential had shrunk from warehouses sometimes the size of seven New York City blocks to a small set of tasks at a workstation no greater than 10 feet by 10 feet.

Each of these examples points to an under-discussed and heterodox aspect of current uses of AI on job quality and skills: current managerial decisions and technological products

mean the transition to the purportedly attainable full automation of a specific job could well be one where the workers in that job experience less autonomy (and thus fewer opportunities for creativity, empathy, complex problem-solving, and judgment), not more.

These "transition to automation" periods can be extremely long, as incremental progress either continues or stalls and researchers work for the next breakthrough. See, for example, the delays relative to predicted timelines for self-driving cars.[78] When thinking about workers training their AI replacements, some may have in mind the time horizon of training another human to do your job — but these periods of automation transition could last years, decades, or possibly the span of an entire career.

Depending on how the technology evolves, workers may never see a paradise of creativity on the other side. This is perhaps a corollary to, or a deepening of, the "paradox of automation's last mile" suggested by Mary Gray and Siddarth Suri in *Ghost Work*.[79] While their work highlighted the possibility that there will always be more work for humans in the quest for full automation, workers' present experience working with AI systems tells us a great deal about what it will look like for many workers to traverse that paradoxical last mile.

This is an issue that comes through more clearly in light of the distribution of skills and tasks throughout the labor market and the ways managers and companies decide to combine human and AI labor. In jobs where companies are actively trying to automate some or all of the tasks, they need workers to produce training data: both for originally building the relevant algorithms and for continuous improvement. The current state of AI chasing the replication of human abilities means automation technologies are largely focused on discrete, narrow tasks — and so, too, are the workers tasked with training them.[80] Given these structural forces, it's not surprising that the ways these AI technologies are deployed reduce workers' scope for exercising these more "human" tasks in their jobs.

If executives and managers see value in using AI technology to free up their workers to perform more "human" or advanced tasks, they cannot assume that any AI tool will meet that goal.

The realities of how current managerial uses of AI technologies transform workers' jobs suggest a need to re-evaluate the optimistic framing in multiple ways. To the extent that executives and managers see value in using AI technology to free up their workers to perform more "human" or advanced tasks, they cannot assume that any AI tool will meet that goal. Nor should developers take for granted that the AI tools they create, as implemented, will free up humans to be more creative or empathetic or to focus on tasks requiring more complex judgment or discernment. Without caution and active collaboration with workers, these workplace AI product adoptions may bring about the very opposite effects. Moreover, the uneven pace of AI development means that these current impacts ought not be brushed aside as temporary harms on a quick path to a better future. The future capabilities of these technologies remain unclear, as do the timelines to achieve them.[81] The present-day harms to existing workers' autonomy, dignity, and senses of satisfying and meaningful work, on the other hand, are real — and accelerating.

THEME 5

Empowering workers early in AI development and implementation increases opportunities to implement AI that benefits workers as well as their employers

The market for workplace AI products is presently structured to address needs and opportunities perceived by company leaders and managers with substantial budgets for AI transformations or integrations. Workers who use these products are multiple layers removed from decision-makers and may also be in different departments or reporting lines. As such, the priorities of AI purchasers are not necessarily those of workers.

Providing workers the opportunity to participate in the creation, design, and implementation of workplace AI is a necessary corrective to approaches that exclude workers only to later require them to use technologies created without their input or the centering of their needs. Not every worker who participated in the research wants these opportunities for input. But comprehensively excluding this group throughout the process or until UX (user experience) or user-testing phases has multiple negative effects.

The data annotators who participated in this research were tasked with helping to improve the ML software they used in their work. They described their team leaders and the developers they worked with as open to suggestions, and they took pride in troubleshooting, bug-spotting, and identifying improvements to the software that were later adopted. Interviewees thought their ideas and suggestions meaningfully improved the tools they worked with. Even in this intentionally participatory environment, however, their reflections revealed some missed opportunities. In broad strokes, they described their role as finding ways to improve the software's effectiveness. Nested underneath this mission were implicit objectives like understanding the nuances of the software's failure modes or identifying improvements for the user interface. Since the software was useful but frustrating when it failed, improvements to its effectiveness also contributed to their own satisfaction.

However, they appeared to consider participation in shaping other aspects of their work or offering ideas for new technologies as outside their responsibilities. Recall, for example, the workers who found it disruptive to their efficiency and flow when the ML software struggled and they had to edit its mistakes. An annotator offered the idea of giving annotators the ability to set those images aside to return to later, allowing them to process all of the successful automations in one batch, and all the tasks requiring their edits in another, rather than bouncing back and forth without a sense of what the next task would require. The suggestion was a process improvement that could improve worker satisfaction and likely task efficiency. When asked in a follow-up question whether the annotator had made that suggestion to their managers or the developers, they responded that their job was to improve the effectiveness of the tool, not ideas like this.

By not opening up the biggest possible spaces for worker participation, or not asking the right questions, designers and implementers are missing productive ideas for new technologies.

What would a worker in this role identify as the biggest issues they'd like tech to help them solve or tasks they'd like assistance with? How could workflows, jobs, and business processes be reimagined for the better from workers' perspectives? What would a welcome technological aid or solution look like from their perspective?

A number of the participating customer service agents, for instance, flagged angry customers as one of the worst parts of their jobs. They had no advance notice of whether their next interaction would be with someone who would be respectful or abusive to them about mistakes by their employer, and which were out of their control. When asked if there were areas where they would welcome AI in their jobs, multiple agents suggested de-escalation technologies, or warning systems so they at least knew what they'd be facing — both of which agents were confident would improve customer experience as well.

Without worker participation from the start, AI developers also lack important information about design and use. What common or uncommon occurrences in the workplace would cause this technology to fail or struggle? What does every experienced worker in this role know that outsiders would find difficult to identify or understand? Moreover, workers left out of the process may be less inclined to trust or adopt AI tools.[82] In an alternative world where workers were included from the start, how much more effective could a given technology be? How much more quickly could it be launched at scale?

Not having the workers who will use the technology "in the room" means that projects get greenlit and products get designs that are, on the whole, not worker-centric. Worker-centricity is one of many possible goals for a product team. Without consistent, empowered advocates for that goal present, it is structurally probable to be deprioritized relative to priorities of senior leaders, designers, and engineers.[83]

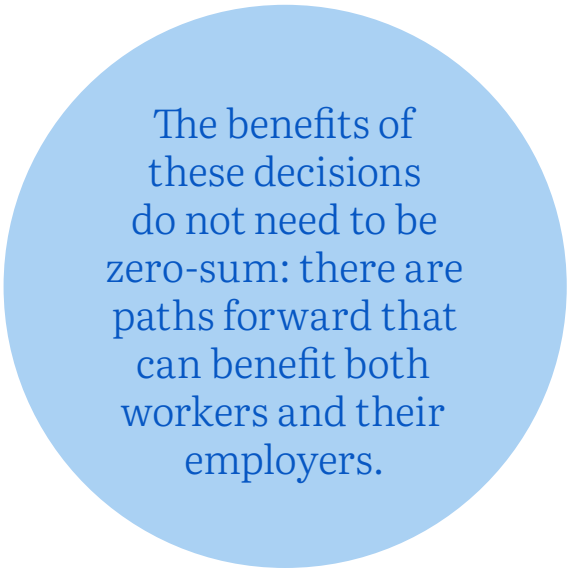
By not asking the right questions, designers and implementers are missing productive ideas for new technologies.

Many workplaceAI systems (including several described in this research) also reflect and reinforce a managerial mindset (perhaps best described as "neo-Fordism" or "neo-Taylorism") where deskilling and strict control over workers is seen as the path to the highest profitability. The origins and drivers of this approach in AI development has not been accounted for in detail (for instance, did limits to AI capabilities shape this approach to workplace AI products, or did strong belief in this managerial approach shape a market that AI developers then filled?), but the impact on workers remains the same — they are treated as subjects in need of discipline and control, rather than as professionals with valuable expertise.

Alternative management approaches used in sectors as varied as manufacturing[84] and hospitality[85] encourage frontline workers to draw on their accumulated expertise and judgment to address problems and make improvements. These approaches afford workers the needed influence and decision rights to make their recommendations stick.[86] [87] Treating workers as genuine experts and empowering them to participate in AI development and deployment offers opportunities for both workers and businesses to benefit.[88]

Opportunities for Impact

Actors across the AI investment, creation, deployment, use, and regulation spectrum have opportunities to make decisions that center workers' voices and protect their well-being. As shown above, the benefits of these decisions do not need to be zero-sum: there are paths forward that can benefit both workers and their employers. Specific recommendations and opportunities for impact are outlined by stakeholder group below. While the recommendations are structured by the audience most able to take action on it, change and accountability also rely on the relationships between different actors (for instance, the complicated ways in which the actions of both AI creators and AI implementers drive the workplace AI products available for purchase). In practice, these interactions may have allowed some decision-makers to evade culpability for their actions; accordingly, these relationships and dynamics should also be considered in the implementation of these recommendations. In the event that actors are attempting to avoid responsibility for harms or negative effects, others must take care to hold them to account.



STAKEHOLDER 1

AI-implementing companies

Employers that choose to use AI in the workplace have an obligation to ensure it does not decrease their employees' well-being. They also have the highest degree of control in ensuring this outcome.[89] While employers might not directly create the AI-enabled workplace products on the market, they can choose which products to use (or choose to use none at all) and set the contexts and conditions for their use.[90] Employers determine when AI is used (e.g., in core or non-core tasks) and how (e.g., as a decision-support tool with a human worker given the ultimate say or as a final decision-making tool). As shown above, this set of decisions has profound influence over how workers experience workplace AI, even in cases where employers are using similar AI products.

OPPORTUNITIES FOR IMPACT

Values and governance	Commit to making worker-centric/worker-friendly AI that increases access to better jobs, especially for the most vulnerable and marginalized workers.
AI product purchasing	Take workers and their institutionalized representatives seriously as experts in their own roles and incorporate their input into purchasing decisions, including: • Which problems and opportunities to seek AI solutions for. (For instance, seeking technology to support workers in their roles in ways that they have identified rather than the current focus on punitive surveillance tools.) • Which solutions to select out of an AI product category.
AI product implementation	Integrate frontline workers and other end-users' perspectives into the implementation of AI (e.g., workflow and performance targets).
	Give humans working directly with AI systems the final judgment on AI-supported decisions, especially in situations where they could affect workers' performance evaluations and lives outside of work.
	Foster and seek out representation from institutionalized forms of worker organization, ensuring that workers can offer their authentic views without fear of retribution or retaliation.

STAKEHOLDER 2

AI-creating companies

Core technologies underlying workplace AI tools are created by an increasingly concentrated group of companies.[91] This concentrated group may use them internally and also sell these technologies to other businesses. Values and practices that center the participation and well-being of worker end-users at these companies have the potential for transformative changes in job quality around the globe. These values and practices are all the more important in the market for workplace AI products, one where company leaders and managers are purchasers and the users may be some of the lowest paid and least influential or powerful employees in the company. This market structure means a focus on customers is not necessarily a focus on worker end-users (and vice versa).

These divergences are likely particularly pronounced in companies with strong command and control approaches to integrating AI into their workplaces as outlined above. Not coincidentally, these companies often employ large pools of low-wage workers most vulnerable to AI's negative effects. Creating better feedback loops and genuinely centering workers will often require seeking out the participation of workers and their representatives beyond their own organizations. There are, however, areas of alignment between the needs and preferences of workers and the incentives of business leaders and managers (as discussed in more detail in Theme 3). While not all of the applications sought by company leaders and managers may be endorsed by their workers, focusing on the overlap adds an additional constituency in support of particular products: the workers/end-users.

OPPORTUNITIES FOR IMPACT

Values and governance	Commit to making worker-centric/worker-friendly AI that increases access to better jobs — especially for the most vulnerable and marginalized workers — by measuring workplace AI products' impacts on job availability, wages, and job quality, and working to eliminate or mitigate negative impacts.
	Include workers as participants and key stakeholders in creating any company's AI ethics/responsible AI principles.[92]
	Recruit staff of diverse backgrounds to AI development teams and actively work to retain them as staff after recruitment.[93] While representation on its own is not a solution, the relative lack of diversity in AI product teams can contribute to the creation of blindspots that could be mitigated by more diverse teams.[94]

AI product origination	Incorporate workers and other end-users' perspectives from the beginning of the product origination process. That is, work forward from problems, challenges, and opportunities identified by frontline and other workers toward products rather than finding ways to shoehorn research progress into workplace products and routines. Collaborate with workers' institutionalized representatives where possible.
	"Red-team" potential use of workplace AI products from origination through major update cycles. Without intentional focus, developer and product teams may not identify the potential for misuse or harm.[95] Eliminate or mitigate identified opportunities for uses harmful to workers, especially in situations where technologies may be sold and deployed in contexts with fewer worker protections than they are developed. Responsible red-teaming and harm mitigation may require companies to not pursue product ideas where harms cannot be mitigated. Particular attention must be paid to the diversity and heterogeneity of use contexts, including ones where potential dimensions of marginalization and inequality (e.g., gender, class, age, ethnicity, race, religion, sexuality, disability status) may not be the same as the cultural and social context of the developing company or team and where existing power imbalances limit the opportunities to reject, restrict, or limit use.[96]
	Collaborate with workers to identify areas where they would welcome assistance in completing their work with augmenting AI or automation of non-core tasks, drawing upon the complementarity of humans and AI.[97]
	Foster and seek out institutionalized representation of workers, ensuring that workers can offer their authentic views without fear of retribution or retaliation.
	When seeking to include the perspectives of workers, recognize that workers from different backgrounds and of different demographic categories may experience workplaces and AI technologies in different ways. Seek broad, representative participation and feedback, and work to ensure workers of all backgrounds feel comfortable and empowered when participating.
AI product development and updating	Take workers seriously as experts in their own roles and include them in product development and future update cycles. Create opportunities for their empowered participation as subject matter experts, not just as end-user testers.

STAKEHOLDER 3

Workers, unions, and worker organizers

Workers and the organizations and unions that represent them can shape AI's impacts on their workplaces through contract negotiations and other mechanisms to influence corporate policy as well as on-going input into purchase and implementation decisions. Unfortunately, it is not common in many countries for employers to invite this participation and the ways AI technology shapes job quality and worker well-being can be obscured. Education and training programs by unions and worker organizations can help workers understand the functions and roles played by AI products and equip workers to participate in decisions made to purchase and implement AI in their workplaces.

OPPORTUNITIES FOR IMPACT

Unions and worker organizations	Foreground worker voice in development and implementation of AI (and other technology) as a plank of contract negotiations and other mechanisms to influence corporate policies.[98]
	Train members and organizers in relevant technologies and their benefits and drawbacks, spotlighting AI technology and related issues (e.g., data rights) as a major influence on working environments.[99]
Workers	Actively seek to participate in workplace AI purchasing and implementation decisions.
	Ask for disclosure and transparency on technologies being used, data being collected, how it's being used, and for what purpose.
	In workplaces with cultures of including workers in management decisions, offer input on areas where AI technology solutions would be welcomed and suggestions about ideal implementation.
	Seek out worker organizations and unions operating in the same sector and geography undertaking efforts on these issues.

STAKEHOLDER 4

Policymakers

Through laws and regulations concerning both technology and labor, government lawmakers and regulators shape the environments in which AI products are developed, sold, and implemented, and thus shape the technologies themselves.[100] As discussed above, there are and will continue to be instances where the incentives of AI-creating and -implementing companies strongly diverge from the interests of their workers. In such instances, government action will be required to ensure the livelihood and well-being of workers; as the historical record indicates, few businesses will voluntarily shoulder the whole of these changes. Compounding this, lack of worker voice and power often comes down to lack of worker protection (e.g., for organizing or ensuring correct worker classification). In some cases, AI technology further enables employers to exploit these power imbalances and policy or enforcement gaps.[101] Strong regulation and enforcement, including of existing laws and policies, is all the more critical in these situations.[102]

Government action will be required to ensure the livelihood and well-being of workers; few businesses will voluntarily shoulder the whole of these changes.

The heavily concentrated nature of the global AI research and workplace product development industry means that many workplace AI technologies are developed in and sold from the United States and China and then implemented in other regulatory environments.[103] While the fractured, global nature of AI's impacts on workers impedes concerted efforts to protect workers, divergent regulatory environments offer opportunities for the experimentation and sharing of best practices in line with local norms and values. Conversely, countries with less economic power or enforcement capacity may find themselves in the position of reacting to harmful technologies created at or implemented from a distance; these situations require careful consideration and differentiated responses.

Much of the African continent, for instance, is both less well-placed to reap the economic benefits of AI (due to a comparative lack of telecom, computing, and other infrastructure, as well as a comparatively small skilled AI workforce), and more susceptible to potential workforce and labor market harms from AI use inside and outside the region (due to a comparative absence of protective regulations targeting AI use and impacts and weaker enforcement capabilities for labor protections). While a number of countries have been making recent strides on these factors, they are beginning from less advantageous starting points and starting at later dates than many high-income countries and regions.[104] Proactively investing in AI workforce development and supporting infrastructure opens up the possibility of more "home grown" solutions responsive to local needs and values, rather than the status quo importation of technology from abroad that may undercut local social goals.[105]

OPPORTUNITIES FOR IMPACT

Worker voice	Safeguard worker organizing on working conditions (e.g., tech introductions and implementations) and unionization through additional legislation and enforcement as needed.
	Give workers the right to know about technologies used in their workplace, the data being collected on them, and the intended uses and impacts of the technology and data.[106]
Worker protection	Where possible, regulate and enforce protections from known harms to workers caused by AI through existing legislation and agencies.
	Create new, targeted legislation and regulations to address gaps in worker protection, either as standalone provisions focused on workers[107] or as a part of broader efforts to regulate AI technologies.[108]
	Protect worker organizing for improved working conditions (e.g., tech introductions and implementations) and unionization.
Tax policy	Identify opportunities to correct the balance of tax burden between labor and capital, which shape the conditions for when and how employers choose to use workplace AI technologies,[109] as well as workers' influence, leverage, or voice in workplaces.
Investment regulation	Require inclusion of relevant worker impact and human capital measurements in standard reporting and disclosure metrics.
Research grants and proposals	Require assessments of anticipated impacts on job availability and quality in government AI research grants.[110]
	Solicit ideas and prototypes of worker-friendly/worker-complementing AI technology (for instance, through RFPs or Grand Challenges) and fund their development with public research and development grants.
Low- and middle-income country (LMIC) responses	Create multi-country and multi-stakeholder collaborations across LMICs facing similar challenges and reform existing multi-stakeholder groups to provide more influence to the least powerful and most vulnerable participating groups. While perspectives from representatives of LMICs are included in some existing global multistakeholder efforts, the structure of these groups and their embedded power imbalances mean participation from less powerful actors may function as a "box ticking" exercise rather than as a true and influential representation of their specific needs.[111] The creation of collaborative groups facing similar challenges would enable them to work together on identifying specific needs, as well as to potentially take collective or coordinated action in addressing them.
	Invest in local AI workforces and infrastructure to support the development of workplace AI technologies that address local needs in line with local values.

STAKEHOLDER 5

Investors

Private investment in AI technologies doubled between 2020 and 2021, directing ever higher amounts to a concentrated group of companies.[112] While angel and venture capital funders have not traditionally focused on the ESG (Environmental, Social, and Governance) impacts of their investments, the push for higher investor responsibility for climate change and sustainability impacts marks a shift in this attitude. Large institutional investors, similarly, are beginning to articulate an investment thesis of "stakeholder capitalism"[113] inclusive of companies' treatment of workers.[114] In the United States, the Securities and Exchange Commission (SEC), which is responsible for regulating government investments, has proposed additional workforce disclosures related to treatment of workers, arguing that they are material information for investors.[115] As AI is increasingly adopted by companies, investors can influence its path by understanding and accounting for the downside risks posed by practices harmful to workers and the potential value created by worker-friendly technologies and practices.

OPPORTUNITIES FOR IMPACT

Investors	Include job availability and quality impacts of AI technology in ESG impact measurements for AI-creating and AI-implementing companies.
	Offer and support shareholder proposals to increase workers' voice and well-being.[116]
	Request anticipated impact on workers when evaluating proposals and pitches for workplace AI products and companies.
	Work with institutionalized forms of worker representation in order to solicit authentic, unencumbered perspectives of workers to incorporate into ESG metrics, stakeholder capitalism initiatives, and other efforts intended to increase worker well-being.

Conclusion

Current uses of AI and its existing technological path pose significant risks to job quality and worker well-being. Future research is needed to better understand the impacts of AI on informal workers, as well as to quantify the business case for and possible tradeoffs of the recommendations in this report. However, opportunities exist for developers, employers, workers and their institutional organizations, governments, and investors to correct the course, steering AI in a direction that benefits workers as well as their employers.

The Partnership on AI is leading an effort — in continued collaboration with workers at the frontier of AI implementation — to develop commitments and targets for AI's impacts on job quality as well as tools to support stakeholders in implementing these targets and commitments. To learn more about this effort and stay updated on our work, visit the AI and Shared Prosperity Initiative page on PAI's website.

The ways AI degrades job quality and worker well-being in the present are neither inherent to the technology nor inevitable outcomes of its progress. Stakeholders have the power to transform AI's trajectory for the better. It is incumbent upon them to use it.

Acknowledgments

This research is based on the reflections and comments of nearly 80 workers in India, sub-Saharan Africa, and the United States who work with AI every day in their jobs. Without the insights and stories they shared in the diary study and their interviews, this research would not have been possible. We thank them for sharing their time and expertise with us.

We appreciate the employees of the company who took an interest in the goals of this research and supported its execution by introducing us to their colleagues whose participation formed the sub-Saharan Africa site, as well as the professional firm which assisted us with recruiting participants at the India and United States sites. Each of them went above and beyond to ensure we were able to connect with people directly experiencing AI's complex effects in their workplaces.

The members of the Steering Committee for Partnership on AI's AI and Shared Prosperity Initiative shared their expertise and feedback on this research throughout the scoping, design, analysis, and writing process. Their astute contributions and detailed comments on earlier drafts have strengthened this work immensely. We are grateful for their guidance. We additionally appreciate Mary L. Gray for taking the time to share her valuable suggestions on early versions of our research questions.

We would also like to thank the Partnership on AI staff who championed this work and provided thoughtful feedback and ideas throughout the research and writing process: Andrea Cross, Christine Custis, Rebecca Finlay, Hudson Hongo, Sonam Jindal, Katya Klinova, Tina Park, and Neil Uhl.

Finally, we would like to thank the Ford Foundation, and the stewardship of Ritse Erumi in particular, as a generous funder of Partnership on AI's AI and Shared Prosperity Initiative and of this research.

The AI and Shared Prosperity Initiative Steering Committee

Daron Acemoglu
Institute Professor of Economics at the Massachusetts Institute of Technology

Dunstan Allison-Hope
Vice President at Business for Social Responsibility

Juana Catalina Becerra Sandoval
Visiting Scholar, Responsible & Inclusive Tech at IBM

Michael Chui
Partner at the McKinsey Global Institute

Jack Clark
Co-chair of the AI Index and former Policy Director at OpenAI

Andrea Dehlendorf
Executive Director at United for Respect

Arturo Franco
Vice President for Insight at the Mastercard Center for Inclusive Growth

Jessica Fulton
Vice President at the Joint Center for Political and Economic Studies

Ryan Gerety
Senior Advisor at United for Respect

Stacey MacGrath Gifford
Research Staff Member—Future of Climate, Governance of Science and Technology at IBM Research

Deborah Greenfield
Senior Advisor at the US Department of Labor and former Deputy Director-General for Policy, International Labour Organization

Anton Korinek
Professor of Economics at the University of Virginia

Andrew Kortina
Co-founder of Fin Analytics and Venmo

Jody Medich
Principal Design Researcher, Office of the CTO at Microsoft and Founder of Superhuman-X

Sabelo Mhlambi
Founder at Bantucracy

Pamela Mishkin
Policy Staff Member at OpenAI

Shakir Mohamed
Senior Staff Scientist at DeepMind

Paloma Muñoz Quick
Advisor at the UN Working Group on Business and Human Rights and OHCHR B-Tech

Grace Mutung'u
Research Fellow at the Centre for Intellectual Property and Information Technology at Strathmore University

Lama Nachman
Intel Fellow & Director of Anticipatory Computing Lab

Reema Nanavati
Executive Director of the Self Employed Women's Association

Aiha Nguyen
Program Director, Labor Futures Initiative at Data & Society

Rahul Panicker
Chief Research and Innovation Officer at the Wadhwani Institute for AI

Damon Silvers
Policy Director and Special Counsel at AFL-CIO

Sarah Treuhaft
Vice President of Research at PolicyLink

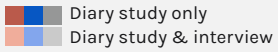
Appendices

APPENDIX 1

Detailed Site and Technology Descriptions

OVERALL

FIG. 3 Participant demographics & characteristics: Overall

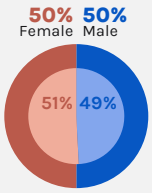


Number of participants

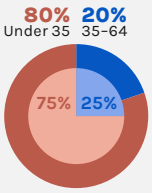
78 Diary studies

47 Interview participants

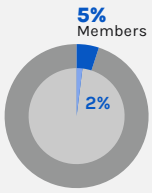
Gender



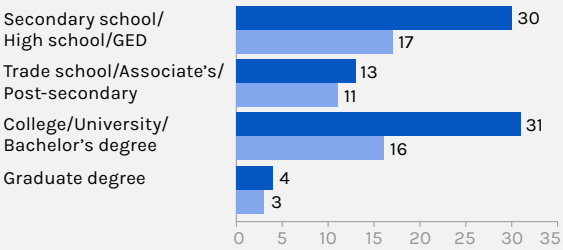
Age



Union membership



Education level completed



Options for gender identification at each site included "Non-binary," "Other (specify)" and "Prefer not to answer." None of the participants chose these options.

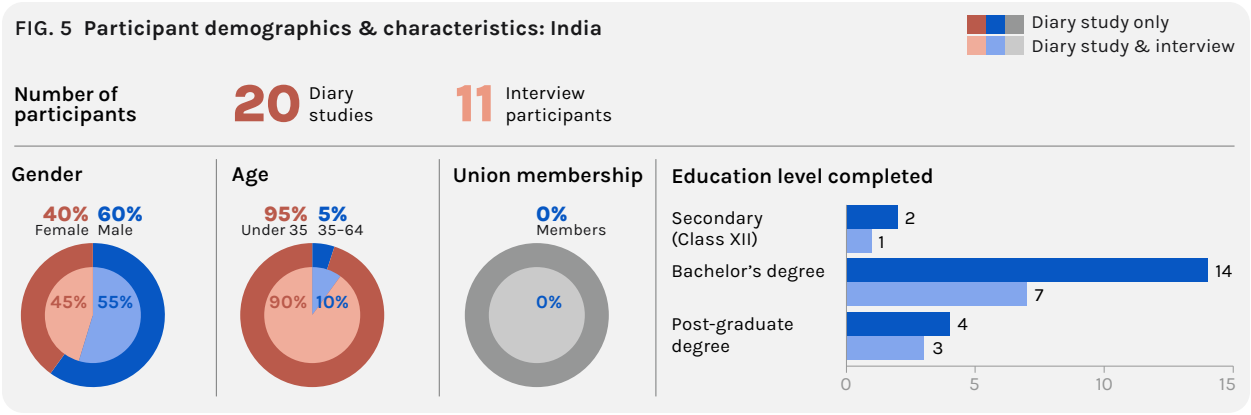
FIG. 4 AI technology types included in study by worker journey stage and site

Employee Journey Stage	Technology Type	India	Sub-Saharan Africa	USA
ROLE FULFILLMENT	Task augmentation	●		
	Task automation	●	●	●
	Real-time task assignment			●
COACHING	Real-time performance feedback	●	●	●
	Performance predictions	●		
	Improvement guidance			●
PERFORMANCE EVALUATION	Performance target setting	●		●
	Performance assessment	●		●

INDIA

ROLE FULFILLMENT	Task augmentation	Software that analyzed text-based customer service chats to provide agent with small set of potential response and resolution templates to use as responses to customer comments
ROLE FULFILLMENT	Task automation	Chatbot that handled straightforward customer inquiries, automating requests that previously would have been handled by workers
		Software that captured and logged customer and call details, which previously agents needed to enter themselves
COACHING	Real-time performance feedback	Software that monitored calls for keywords, key phrases, and tone, and provided pop-up coaching to the call agent in the moment on what to do differently (e.g., greet the customer according to the script; speak more loudly, quietly, slowly, quickly; end the call by a certain time)
		Software that monitored calls and predicted customer satisfaction scores
PERFORMANCE ASSESSMENT	Performance evaluation	Software that monitored calls and predicted customer satisfaction scores

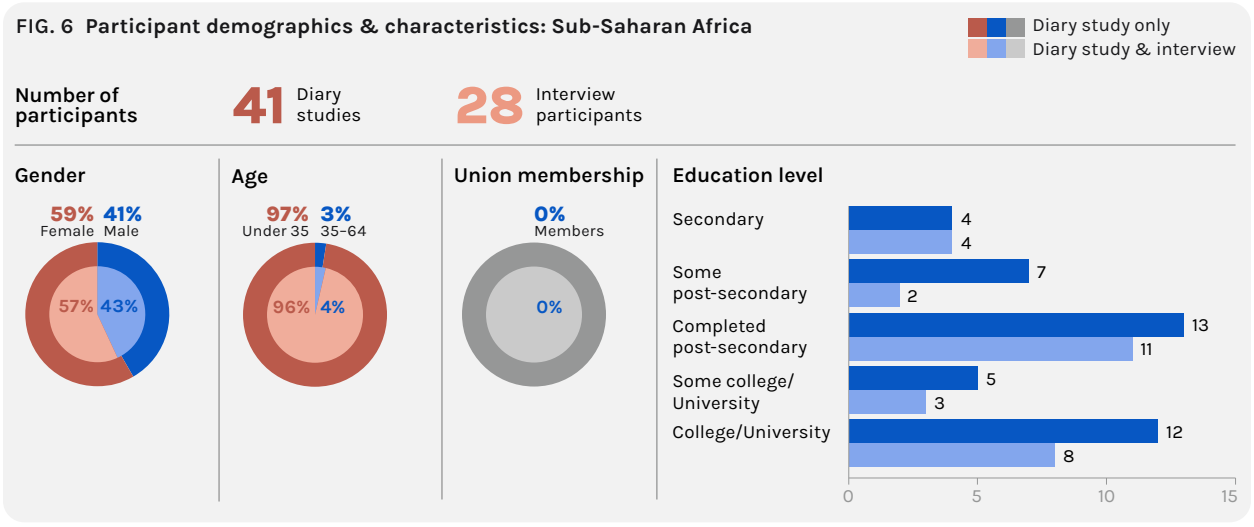
FIG. 5 Participant demographics & characteristics: India



SUB-SAHARAN AFRICA

ROLE FULFILLMENT	Task automation	Software used in annotating videos to provide AI model training data; enables workers to annotate an object in one frame, and have the annotation of that object continue through many frames. Previously workers had to annotate the object in each frame of the video
		Software used in annotating images, allows workers to select the far corners of the object that needs to be annotated, and the tool will draw the rest of the contours around the object
COACHING	Real-time performance feedback	Software used to review the accuracy of annotations. Approves or rejects, does not provide any kind of feedback or coaching on what needs to be improved

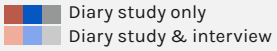
FIG. 6 Participant demographics & characteristics: Sub-Saharan Africa

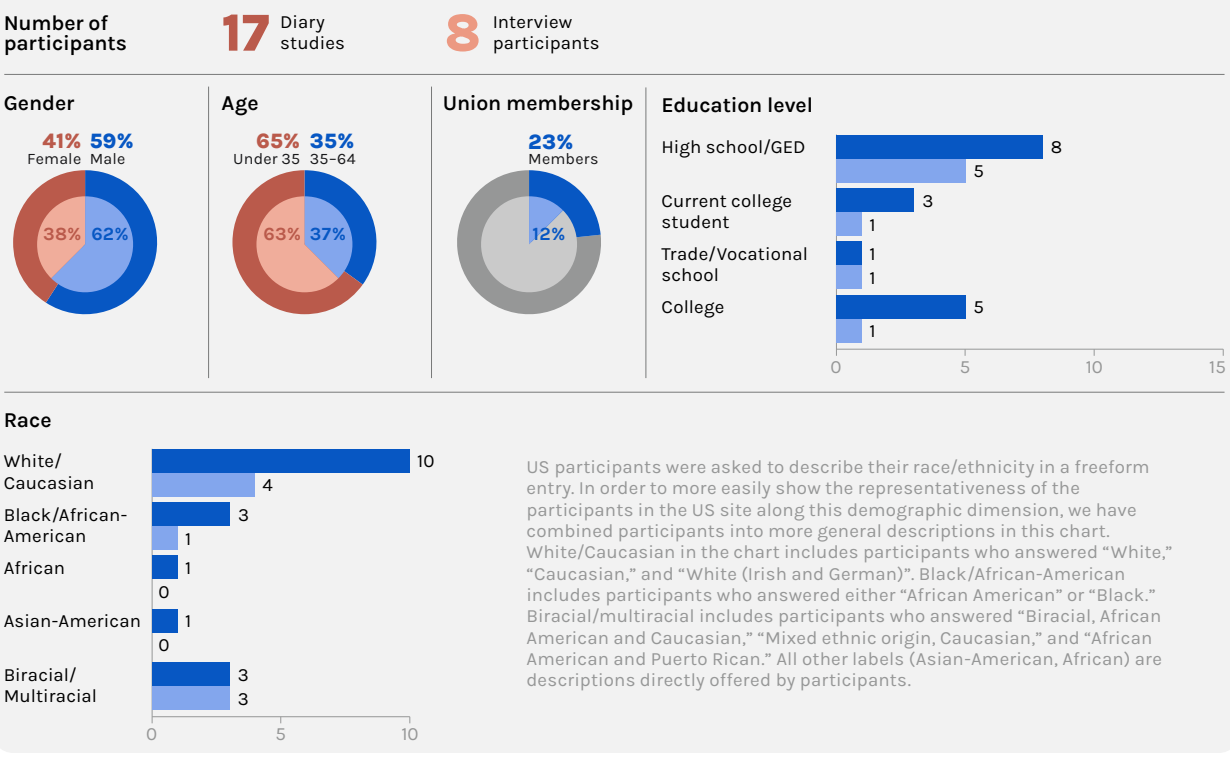


UNITED STATES

ROLE FULFILLMENT	Real-time task assignment	Handheld scanner task assignment: Device and software used to complete tasks (e.g., give workers instructions to guide the movement of items throughout the warehouse, or which trucks to load packages onto)
ROLE FULFILLMENT	Task augmentation	Return processing hardware and software: Device and software to assess condition of returns, and direct worker on next steps for item (e.g., repackage, dispose)
		Robotic item movers: Robots deliver items to pickers, packers, and stowers, reducing or eliminating the need for humans to move around the warehouse to complete these jobs or carry heavy loads
COACHING	Real-time performance feedback	Handheld scanner feedback: Device and software used to complete tasks also provides information on: • Whether tasks have been correctly completed (e.g., the worker has found the correct item, placed the item in the correct location, or loaded the package onto the correct truck) • The worker's speed of task completion (how long between completing one task and the next), and number of tasks completed per hour
PERFORMANCE EVALUATION	Performance target setting	Software that sets targets for worksite based on historical performance and company goals
	Performance assessment	Software tracks performance against targets, warns worker and notifies management when performance lags for too long

FIG. 7 Participant demographics & characteristics: United States





US participants were asked to describe their race/ethnicity in a freeform entry. In order to more easily show the representativeness of the participants in the US site along this demographic dimension, we have combined participants into more general descriptions in this chart. White/Caucasian in the chart includes participants who answered "White," "Caucasian," and "White (Irish and German)". Black/African-American includes participants who answered either "African American" or "Black." Biracial/multiracial includes participants who answered "Biracial, African American and Caucasian," "Mixed ethnic origin, Caucasian," and "African American and Puerto Rican." All other labels (Asian-American, African) are descriptions directly offered by participants.

APPENDIX 2

Research Methodology

Research methods

The themes and recommendations from this report are grounded in a literature review of related subjects, as well as primary research with workers at each site, virtually fielded from November 2021 until March 2022. All research was conducted in English, unless noted below. Where the literature revealed opportunities to better understand AI's impact on job quality and worker well-being, we incorporated these areas into guides for diary studies to be completed by participating workers. These diary studies then served as a foundation for a series of semi-structured interviews with diary study participants at each site.

Participants shared their reflections on the diary study questions and prompts through a combination of text, voice, and video answers, as well as image responses to a limited number of questions. To better understand themes emerging from the diary studies, a subset of the diary study participants were each invited to participate in a 60-minute semi-structured interview via Zoom where they were asked a combination of universal questions about emerging themes as well as specific questions about their diary study and interview responses. All research was conducted in English except where noted below. The participants had high degrees of competency in English due to living in countries where English is frequently spoken, using it as the main language in their jobs, or both. For the participants based in India, these interviews were conducted one-on-one in English with the primary researcher. A translator fluent in Hindi was also on the call to offer translation support if requested by participants. US interviewees participated in one-on-one Zoom calls with the primary researcher.

Unlike the other sites, participants in the sub-Saharan Africa site were colleagues. To prevent the possible appearance of choosing favorites by selecting only a subset of participants for interviews, all participants at this site who completed the diary study were invited to participate in interviews. Due to strong interest from the participants who completed the diary study and comparatively limited research team time, these interviews were structured as group interviews with three to four participants per group. Due to last-minute schedule changes from participants, the actual number of interviewees per group ranged from one to four. The research team explored possible upsides and downsides of this approach with leads from the sub-Saharan African group, who offered guidance that this approach should work well; employees were used to participating in focus groups and other similar discussions with each other. This participant site included several individuals with team/project management, mentorship, or coaching responsibilities. Each of these individuals was interviewed one-on-one to ensure that they could speak freely about aspects of their managerial, supervision, or mentorship duties they might prefer not to discuss in front of potential reports. Similarly, this design ensured that workers could speak about experiences they might prefer not to share or discuss with managers or supervisors.

Participant recruitment

For this work, the research team sought to include a wide range of possible workers who frequently use or interact with AI in their workplaces with a focus on workers who were structurally likely to be more vulnerable to any harms from these technologies. We sought to learn from:

- Vulnerable workers in a high income country (e.g., workers with fewer formal education credentials and thus fewer opportunities for higher paying jobs)
- Middle class workers in a LMIC (e.g., workers who were less vulnerable in the context of their own country, but could be substantively affected by global market changes driven by decisions abroad or in high income countries)
- Working class or working poor workers in a LMIC (e.g., workers who could be both individually and collectively affected by the forces described for the prior two sites, and thus at highest risk of harm).

The team considered three possible approaches for recruitment:

Recruitment Approach	Advantages	Disadvantages and Mitigations
Working with an independent research recruiter	Workers could participate independently of their employer	Participants would need their anonymity protected to ensure their employers would not punish their participation Lack of employer permission might make participants feel less comfortable speaking frankly about their experiences
Collaborating with a supportive employer	Workers participate with their employer's full permission	Participants' individual and collective ability to participate voluntarily and speak freely would need to be protected through confidentiality and anonymity provisions
Collaborating with a worker organization or union	Workers could participate independently of their employer Workers would likely have higher familiarity with how technologies affect working conditions and more	Employers often scrutinize organizers and active union participants more heavily, increasing their possible risk of participation and the importance of protecting their anonymity High levels of familiarity with technology's impacts built through organizing work might not be representative of the broader group of workers Worker organizations/unions are rare or non-existent in the LMIC occupations that work closely with AI technologies

Due to the nature of the research, there was no ideal set of sites to conduct this work. Each approach would require: (a) setting up provisions to protect participants' ability to voluntarily participate and speak freely, and (b) attention by the research team to ensure the representativeness of the participant pools and the perspectives they shared. The team ultimately decided against recruiting through worker organizations and unions, due to the difficulty of this approach in the LMIC site, and concerns about representativeness in the high-income country.

For the participants based in India and the United States, we worked with a professional research recruiter to identify and invite qualified participants for the study. Participants were recruited through advertisements and targeted outreach.

For the participants based in sub-Saharan Africa, we collaborated with a company which was developing a series of ML tools to assist their data annotators in completing their work. This company agreed to identify a group of employees who were experienced in using these tools and to forward them an introductory note from the research team which outlined the research and invited participation. Participation was entirely voluntary — not required or encouraged by the company, as was made clear in the introductory note. If participants were interested, they were asked to sign up and participate on their own time, using their personal phones or devices (rather than during their work hours, using company devices). The note was forwarded to the group's personal email addresses (not their work addresses) to further underscore the independence of the research and voluntary nature of participation.

All participants at all sites were compensated for their time, with interview participants receiving an additional amount for their additional time. Compensation amounts were set to be generous relative to participants' normal hourly wages, without being so high as to create undue pressure on participants to join the study (and thus potentially undermine the voluntariness of their participation).

Ethics and informed consent

In line with best practices in qualitative research, each participant in the study was informed of the goals, content and format of the study, the benefits and risks of their participation, and the organization and individuals responsible for the study. They were additionally informed that any stories or quotes shared in public research outputs would be anonymized to protect their identity and mitigate any risks of their participation. All participants digitally signed consent forms confirming that they received and understood the information about the study and that their participation was entirely voluntary and could be withdrawn by them for any reason at any time. They were also reminded of this at the outset of each interview, where they were also informed that they could choose to decline to answer any questions posed to them and choose to end the interview at any time.

Since the introduction to the participants at the sub-Saharan Africa site was performed

by participants' employer, additional information was included in the call for participants and the consent form for that site to clarify that the research was being conducted and managed by an external group. The call for participation also included the confidentiality protections participants could expect, namely that their individual responses would not be shared with their employer. The participating company additionally signed a memorandum of understanding in advance of embarking on the study confirming that: (a) no individual or attributable data from participants would be shared with them, and (b) any information shared with them from the research would be synthesized and shared in the form of high-level themes.

For each site, the research team at PAI was the only group with access to participants' diary study and interview responses. In the case of group interviews, participants agreed at the start of interviews to keep any comments shared by others in the group confidential and to protect the privacy of other participants. In line with standard qualitative research practices, the employers at all sites are not named, to assist in protecting the anonymity of participants.

Confidentiality and data storage

All participants' identities have been anonymized in the research output. After completion of the diary studies and interviews, participant responses have been stored separately from identifying information about the participants who provided them.

Data analysis

Information from this study has been analyzed using two approaches. To determine a framework of workplace AI product types from workers' perspectives and to categorize the technologies mentioned in this study according to the framework, a deductive approach was used. To identify themes from the participant responses, an iterative inductive approach was used both to identify initial and emerging themes and to synthesize and cluster those initial themes into the major themes in this report.

References

1 Daniel Zhang et al., "The AI Index 2022 Annual Report" (AI Index Steering Committee, Stanford Institute for Human-Centered AI, Stanford University, March 2022), https://aiindex.stanford.edu/wp-content/uploads/2022/03/2022-AI-Index-Report_Master.pdf.

2 Michael Chui et al., "Global AI Survey 2021," Survey (McKinsey & Company, December 8, 2021), http://ceros.mckinsey.com/global-ai-survey-2020-a-desktop-3-1.

3 Jacques Bughin et al., "Artificial Intelligence: The Next Digital Frontier?," Discussion Paper (McKinsey Global Institute, June 2017), https://www.mckinsey.com/~/media/mckinsey/industries/advanced%20electronics/our%20insights/how%20artificial%20intelligence%20can%20deliver%20real%20value%20to%20companies/mgi-artificial-intelligence-discussion-paper.ashx.

4 Partnership on AI, "Redesigning AI for Shared Prosperity: An Agenda" (Partnership on AI, May 2021), https://partnershiponai.org/paper/redesigning-ai-agenda/.

5 David Autor, David A. Mindell, and Elisabeth B. Reynolds, The Work of the Future: Building Better Jobs in an Age of Intelligent Machines (The MIT Press, 2022), https://doi.org/10.7551/mitpress/14109.001.0001.

6 Daniel Zhang et al., "The AI Index 2022 Annual Report" (AI Index Steering Committee, Stanford Institute for Human-Centered AI, Stanford University, March 2022), https://aiindex.stanford.edu/wp-content/uploads/2022/03/2022-AI-Index-Report_Master.pdf.

7 Lant Pritchett, "The Future of Jobs Is Facing One, Maybe Two, of the Biggest Price Distortions Ever," Middle East Development Journal 12, no. 1 (January 2, 2020): 131–56, https://doi.org/10.1080/17938120.2020.1714347.

8 James K. Harter, Frank L. Schmidt, and Theodore L. Hayes, "Business-Unit-Level Relationship between Employee Satisfaction, Employee Engagement, and Business Outcomes: A Meta-Analysis," Journal of Applied Psychology 87, no. 2 (2002): 268–79, https://doi.org/10.1037/0021-9010.87.2.268.

9 Kaoru Ishikawa, What Is Total Quality Control? The Japanese Way, trans. David John Lu (Englewood Cliffs, N.J.: Prentice-Hall, 1985).

10 Gary P. Pisano, The Development Factory: Unlocking the Potential of Process Innovation (Harvard Business Press, 1997).

11 Terje Slåtten and Mehmet Mehmetoglu, "Antecedents and Effects of Engaged Frontline Employees: A Study from the Hospitality Industry," in New Perspectives in Employee Engagement in Human Resources (Emerald Group Publishing, 2015).

12 Kayhan Tajeddini, Emma Martin, and Levent Altinay, "The Importance of Human-Related Factors on Service Innovation and Performance," International Journal of Hospitality Management 85 (February 1, 2020): 102431, https://doi.org/10.1016/j.ijhm.2019.102431.

13 Sergio Fernandez and David W. Pitts, "Understanding Employee Motivation to Innovate: Evidence from Front Line Employees in United States Federal Agencies," Australian Journal of Public Administration 70, no. 2 (2011): 202–22, https://doi.org/10.1111/j.1467-8500.2011.00726.x.

14 Edward P. Lazear, "Compensation and Incentives in the Workplace," Journal of Economic Perspectives 32, no. 3 (August 2018): 195–214, https://doi.org/10.1257/jep.32.3.195.

15 Joan Robinson, The Economics of Imperfect Competition (Springer, 1969).

16 José Azar, Ioana Marinescu, and Marshall I. Steinbaum, "Labor Market Concentration," Working Paper, Working Paper Series (National Bureau of Economic Research, December 2017), https://doi.org/10.3386/w24147.

17 Alan Manning, Monopsony in Motion: Imperfect Competition in Labor Markets, Monopsony in Motion (Princeton University Press, 2013), https://doi.org/10.1515/9781400850679.

18 Caitlin Lustig et al., "Algorithmic Authority: The Ethics, Politics, and Economics of Algorithms That Interpret, Decide, and Manage," in Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems, CHI EA '16 (New York, NY, USA: Association for Computing Machinery, 2016), 1057–62, https://doi.org/10.1145/2851581.2886426.

19 Aiha Nguyen, "The Constant Boss: Work Under Digital Surveillance" (Data and Society, May 2021), https://datasociety.net/library/the-constant-boss/.

20 Matt Scherer, "Warning: Bossware May Be Hazardous to Your Health" (Center for Democracy & Technology, July 2021), https://cdt.org/wp-content/uploads/2021/07/2021-07-29-Warning-Bossware-May-Be-Hazardous-To-Your-Health-Final.pdf.

21 Mary L. Gray and Siddharth Suri, Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass (Houghton Mifflin Harcourt, 2019).

22 Alexandra Mateescu and Aiha Nguyen, "Algorithmic Management in the Workplace," Explainer (Data and Society, February 2019), https://datasociety.net/wp-content/uploads/2019/02/DS_Algorithmic_Management_Explainer.pdf.

23 Daniel Schneider and Kristen Harknett, "Schedule Instability and Unpredictability and Worker and Family Health and Wellbeing," Working Paper (Washington Center for Equitable Growth, September 2016), http://cdn.equitablegrowth.org/wp-content/uploads/2016/09/12135618/091216-WP-Schedule-instability-and-unpredictability.pdf.

24 V.B. Dubal. "Wage Slave or Entrepreneur?: Contesting the Dualism of Legal Worker Identities." California Law Review 105, no. 1 (2017): 65–123, https://www.jstor.org/stable/24915689.

25 Ramiro Albrieu, ed., Cracking the Future of Work: Automation and Labor Platforms in the Global South, 2021, https://fowigs.net/publication/cracking-the-future-of-work-automation-and-laborplatforms-in-the-global-south/.

26 Phoebe V. Moore, "OSH and the Future of Work: Benefits and Risks of Artificial Intelligence Tools in Workplaces," Discussion Paper (European Agency for Safety and Health at Work, 2019), https://osha.europa.eu/en/publications/osh-and-future-work-benefits-and-risks-artificial-intelligence-tools-workplaces.

27 Frank Pasquale, The Black Box Society: The Secret Algorithms That Control Money and Information (Harvard University Press, 2015).

28 Ifeoma Ajunwa, "The 'Black Box' at Work," Big Data & Society 7, no. 2 (July 1, 2020): 2053951720966181, https://doi.org/10.1177/2053951720938093.

29 Isabel Ebert, Isabelle Wildhaber, and Jeremias Adams-Prassl, "Big Data in the Workplace: Privacy Due Diligence as a Human Rights-Based Approach to Employee Privacy Protection," Big Data & Society 8, no. 1 (January 1, 2021): 20539517211013052, https://doi.org/10.1177/20539517211013051.

30 Andrea Dehlendorf and Ryan Gerety, "The Punitive Potential of AI," in Redesigning AI, Boston Review (MIT Press, 2021), https://bostonreview.net/forum_response/the-punitive-potential-of-ai/.

31 Partnership on AI, "Framework for Promoting Workforce Well-Being in the AI-Integrated Workplace" (Partnership on AI, August 2020), https://partnershiponai.org/paper/workforce-wellbeing/.

32 This report does not go into detail on financial well-being. Partnership on AI's AI and Shared Prosperity Initiative and team addresses this important component of worker well-being in other publications, including Katya Klinova and Anton Korinek, "AI and Shared Prosperity," in Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society, 2021, 645–51, https://doi.org/10.1145/3461702.3462619, Ekaterina Klinova, "Governing AI to Advance Shared Prosperity," in The Oxford Handbook of AI Governance (Oxford University Press, 2022), https://doi.org/10.1093/oxfordhb/9780197579329.013.43, and Anton Korinek and Megan Juelfs, "Preparing for the (Non-Existent?) Future of Work," NBER Working Paper Series, June 2022, 43.

33 Karen Hao, "Artificial Intelligence Is Creating a New Colonial World Order," MIT Technology Review, accessed July 24, 2022, https://www.technologyreview.com/2022/04/19/1049592/artificial-intelligence-colonialism/.

34 Shakir Mohamed, Marie-Therese Png, and William Isaac, "Decolonial AI: Decolonial Theory as Sociotechnical Foresight in Artificial Intelligence," Philosophy & Technology 33 (December 1, 2020), https://doi.org/10.1007/s13347-020-00405-8.

35 Aarathi Krishnan et al., "Decolonial AI Manyfesto," accessed July 24, 2022, https://manyfesto.ai/.

36 OECD.AI (2021), powered by EC/OECD (2021). "Database of National AI Policies." https://oecd.ai/en/dashboards.

37 Kofi Yeboah, "Artificial Intelligence in Sub-Saharan Africa: Ensuring Inclusivity." (Paradigm Initiative, December 2021), https://paradigmhq.org/report/artificial-intelligence-in-sub-saharan-africa-ensuring-inclusivity/.

38 Adapted from Qualtrics' employee lifecycle model, "Employee Lifecycle: The 7 Stages Every Employer Must Understand and Improve," Qualtrics, https://www.qualtrics.com/experience-management/employee/employee-lifecycle/.

39 Ifeoma Ajunwa and Daniel Greene, "Platforms at Work: Automated Hiring Platforms and Other New Intermediaries in the Organization of Work," in Work and Labor in the Digital Age, ed. Steve P. Vallas and Anne Kovalainen, vol. 33, Research in the Sociology of Work (Emerald Publishing Limited, 2019), 61–91, https://doi.org/10.1108/S0277-283320190000033005.

40 Miranda Bogen, "All the Ways Hiring Algorithms Can Introduce Bias," Harvard Business Review, May 6, 2019, https://hbr.org/2019/05/all-the-ways-hiring-algorithms-can-introduce-bias.

41 Data and Trust Alliance, "Algorithmic Bias Safeguards for Workforce," January 2022, https://dataandtrustalliance.org/Algorithmic_Bias_Safeguards_for_Workforce_Overview.pdf.

42 Mona Sloane, Emanuel Moss, and Rumman Chowdhury, "A Silicon Valley Love Triangle: Hiring Algorithms, Pseudo-Science, and the Quest for Auditability," Patterns 3, no. 2 (February 11, 2022): 100425, https://doi.org/10.1016/j.patter.2021.100425.

43 U.S. Bureau of Labor Statistics. "All Employees, Warehousing and Storage." FRED, Federal Reserve Bank of St. Louis. FRED, Federal Reserve Bank of St. Louis, July 2022. https://fred.stlouisfed.org/series/CES4349300001.

44 Lee Rainie et al., "AI and Human Enhancement: Americans' Openness Is Tempered by a Range of Concerns" (Pew Research Center, March 2022), https://www.pewresearch.org/internet/wp-content/uploads/sites/9/2022/03/PS_2022.03.17_AI-HE_REPORT.pdf.

45 James Manyika et al., "Jobs Lost, Jobs Gained: What the Future of Work Will Mean for Jobs, Skills, and Wages" (McKinsey Global Institute, November 28, 2017), https://www.mckinsey.com/featured-insights/future-of-work/jobs-lost-jobs-gained-what-the-future-of-work-will-mean-for-jobs-skills-and-wages.

46 Mary L. Gray and Siddharth Suri, Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass (Houghton Mifflin Harcourt, 2019).

47 Mayank Kumar Golpelwar, Global Call Center Employees in India: Work and Life between Globalization and Tradition (Springer, 2015).

48 Hye Jin Rho, Shawn Fremstad, and Hayley Brown, "A Basic Demographic Profile of Workers in Frontline Industries" (Center for Economic and Policy Research, April 2020), https://cepr.net/wp-content/uploads/2020/04/2020-04-Frontline-Workers.pdf.

49 International Labour Office. “Women and Men in the Informal Economy: A Statistical Picture (Third Edition).” International Labour Office, 2018. http://www.ilo.org/wcmsp5/groups/public/---dgreports/---dcomm/documents/publication/wcms_626831.pdf.

50 International Labour Office. “Women and Men in the Informal Economy: A Statistical Picture (Third Edition).” International Labour Office, 2018. http://www.ilo.org/wcmsp5/groups/public/---dgreports/---dcomm/documents/publication/wcms_626831.pdf.

51 OECD, and International Labour Organization. “Tackling Vulnerability in the Informal Economy,” 2019. https://www.oecd-ilibrary.org/content/publication/939b7bcd-en.

52 James C. Scott, Seeing like a State: How Certain Schemes to Improve the Human Condition Have Failed, Yale Agrarian Studies (New Haven, Conn.: Yale Univ. Press, 2008).

53 Reema Nanavaty, Expert interview with Reema Nanavaty, Director of Self Employed Women’s Association (SEWA), July 11, 2022.

54 Paul E. Spector, “Perceived Control by Employees: A Meta-Analysis of Studies Concerning Autonomy and Participation at Work,” Human Relations 39, no. 11 (November 1, 1986): 1005–16, https://doi.org/10.1177/001872678603901104.

55 Henry Ongori, “A Review of the Literature on Employee Turnover,” African Journal of Business Management 1, no. 3 (June 30, 2007): 049–054, https://doi.org/10.5897/AJBM.9000028.

56 See Virginia Doellgast and Sean O’Brady, “Making Call Center Jobs Better: The Relationship between Management Practices and Worker Stress,” June 1, 2020, https://ecommons.cornell.edu/handle/1813/74307. for additional detail and impacts of punitive managerial uses of monitoring technology in call centers, including increased worker stress.

57 Aiha Nguyen, “The Constant Boss: Work Under Digital Surveillance” (Data and Society, May 2021), https://datasociety.net/library/the-constant-boss/.

58 Matt Scherer, “Warning: Bossware May Be Hazardous to Your Health” (Center for Democracy & Technology, July 2021), https://cdt.org/wp-content/uploads/2021/07/2021-07-29-Warning-Bossware-May-Be-Hazardous-To-Your-Health-Final.pdf.

59 Mary L. Gray and Siddharth Suri, Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass (Houghton Mifflin Harcourt, 2019).

60 Alexandra Mateescu and Aiha Nguyen, “Algorithmic Management in the Workplace,” Explainer (Data and Society, February 2019), https://datasociety.net/wp-content/uploads/2019/02/DS_Algorithmic_Management_Explainer.pdf.

61 Andrea Dehlendorf and Ryan Gerety, “The Punitive Potential of AI,” in Redesigning AI, Boston Review (MIT Press, 2021), https://bostonreview.net/forum_response/the-punitive-potential-of-ai/.

62 Human Impact Partners and Warehouse Worker Resource Center, “The Public Health Crisis Hidden in Amazon Warehouses,” January 2021, https://humanimpact.org/wp-content/uploads/2021/01/The-Public-Health-Crisis-Hidden-In-Amazon-Warehouses-HIP-WWRC-01-21.pdf.

63 V.B. Dubal. “Wage Slave or Entrepreneur?: Contesting the Dualism of Legal Worker Identities.” California Law Review 105, no. 1 (2017): 65–123, https://www.jstor.org/stable/24915689.

64 Ramiro Albrieu, ed., Cracking the Future of Work: Automation and Labor Platforms in the Global South, 2021, https://fowigs.net/publication/cracking-the-future-of-work-automation-and-laborplatforms-in-the-global-south/.

65 Daniel Schneider and Kristen Harknett, “Schedule Instability and Unpredictability and Worker and Family Health and Wellbeing,” Working Paper (Washington Center for Equitable Growth, September 2016), http://cdn.equitablegrowth.org/wp-content/uploads/2016/09/12135618/091216-WP-Schedule-instability-and-unpredictability.pdf.

66 Arvind Narayanan, “How to Recognize AI Snake Oil,” https://www.cs.princeton.edu/~arvindn/talks/MIT-STS-AI-snakeoil.pdf.

67 Frederike Kaltheuner, ed., Fake AI (Meatspace Press, 2021), https://fakeaibook.com.

68 Aiha Nguyen, “The Constant Boss: Work Under Digital Surveillance” (Data and Society, May 2021), https://datasociety.net/library/the-constant-boss/.

69 Strategic Organizing Center, “Primed for Pain,” May 2021, https://thesoc.org/wp-content/uploads/2021/02/PrimedForPain.pdf.

70 Alessandro Delfanti and Bronwyn Frey, “Humanly Extended Automation or the Future of Work Seen through Amazon Patents,” Science, Technology, & Human Values 46, no. 3 (May 1, 2021): 655–82, https://doi.org/10.1177/0162243920943665.

71 Phoebe V. Moore, “OSH and the Future of Work: Benefits and Risks of Artificial Intelligence Tools in Workplaces,” Discussion Paper (European Agency for Safety and Health at Work, 2019), https://osha.europa.eu/en/publications/osh-and-future-work-benefits-and-risks-artificial-intelligence-tools-workplaces.

72 Strategic Organizing Center, “Primed for Pain,” May 2021, https://thesoc.org/wp-content/uploads/2021/02/PrimedForPain.pdf.

73 Annette Bernhardt, Lisa Kresge, and Reem Suleiman, “Data and Algorithms at Work: The Case for Worker” (UC Berkeley Labor Center, November 2021), https://laborcenter.berkeley.edu/wp-content/uploads/2021/11/Data-and-Algorithms-at-Work.pdf.

74 Andrea Dehlendorf and Ryan Gerety, “The Punitive Potential of AI,” in Redesigning AI, Boston Review (MIT Press, 2021), https://bostonreview.net/forum_response/the-punitive-potential-of-ai/.

75 Carl Benedikt Frey and Michael A. Osborne, "The Future of Employment: How Susceptible Are Jobs to Computerisation?," Technological Forecasting and Social Change 114 (January 1, 2017): 254–80, https://doi.org/10.1016/j.techfore.2016.08.019.

76 "These Are the Top 10 Job Skills of Tomorrow – and How Long It Takes to Learn Them," World Economic Forum, https://www.weforum.org/agenda/2020/10/top-10-work-skills-of-tomorrow-how-long-it-takes-to-learn-them/.

77 Daniel Susskind, "Technological Unemployment," in The Oxford Handbook of AI Governance, ed. Justin Bullock et al. (Oxford University Press), https://doi.org/10.1093/oxfordhb/9780197579329.013.42.

78 Christopher Mims, "Self-Driving Cars Could Be Decades Away, No Matter What Elon Musk Said," WSJ, https://www.wsj.com/articles/self-driving-cars-could-be-decades-away-no-matter-what-elon-musk-said-11622865615.

79 Mary L. Gray and Siddharth Suri, Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass (Houghton Mifflin Harcourt, 2019).

80 Erik Brynjolfsson, "The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence," January 11, 2022, https://doi.org/10.48550/arXiv.2201.04200.

81 World Economic Forum. "Positive AI Economic Futures." Insight Report. World Economic Forum, November 2021. https://www.weforum.org/reports/positive-ai-economic-futures/.

82 Nithya Sambasivan and Rajesh Veeraraghavan, "The Deskilling of Domain Expertise in AI Development," in Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems, CHI '22 (New York, NY, USA: Association for Computing Machinery, 2022), 1–14, https://doi.org/10.1145/3491102.3517578.

83 Sabrina Genz, Lutz Bellmann, and Britta Matthes, "Do German Works Councils Counter or Foster the Implementation of Digital Technologies?," Jahrbücher Für Nationalökonomie Und Statistik 239, no. 3 (June 1, 2019): 523–64, https://doi.org/10.1515/jbnst-2017-0160.

84 Jeffrey K. Liker, The Toyota Way: 14 Management Principles From the World's Greatest Manufacturer (McGraw Hill Professional, 2003).

85 Taiichi Ohno, Toyota Production System: Beyond Large-Scale Production (CRC Press, 1988).

86 Alan G. Robinson and Dean M. Schroeder, "The Role of Front-Line Ideas in Lean Performance Improvement," Quality Management Journal 16, no. 4 (January 1, 2009): 27–40, https://doi.org/10.1080/10686967.2009.11918248.

87 Kayhan Tajeddini, Emma Martin, and Levent Altinay, "The Importance of Human-Related Factors on Service Innovation and Performance," International Journal of Hospitality Management 85 (February 1, 2020): 102431, https://doi.org/10.1016/j.ijhm.2019.102431.

88 Katherine C. Kellogg, Mark Sendak, and Suresh Balu, "AI on the Front Lines," MIT Sloan Management Review, May 4, 2022, https://sloanreview.mit.edu/article/ai-on-the-front-lines/.

89 Zeynep Ton, "The Good Jobs Solution," Harvard Business Review, 2017, 32. https://goodjobsinstitute.org/wp-content/uploads/2018/03/Good-Jobs-Solution-Full-Report.pdf.

90 Abigail Gilbert et al., "Case for Importance: Understanding the Impacts of Technology Adoption on 'Good Work'" (Institute for the Future of Work, May 2022), https://uploads-ssl.webflow.com/5f57d40eb1c2ef22d8a8ca7e/62a72d3439edd66ed6f79654_IFOW_Case%20for%20Importance.pdf.

91 Daniel Zhang et al., "The AI Index 2022 Annual Report" (AI Index Steering Committee, Stanford Institute for Human-Centered AI, Stanford University, March 2022), https://aiindex.stanford.edu/wp-content/uploads/2022/03/2022-AI-Index-Report_Master.pdf.

92 Julian Posada, "The Future of Work Is Here: Toward a Comprehensive Approach to Artificial Intelligence and Labour," Ethics of AI in Context, 2020, http://arxiv.org/abs/2007.05843.

93 Jeffrey Brown, "The Role of Attrition in AI's 'Diversity Problem'" (Partnership on AI, April 2021), https://partnershiponai.org//wp-content/uploads/dlm_uploads/2022/04/PAI_researchpaper_aftertheoffer.pdf.

94 Tina M. Park, "Making AI Inclusive: 4 Guiding Principles for Ethical Engagement" (Partnership on AI, July 2022), https://partnershiponai.org//wp-content/uploads/dlm_uploads/2022/07/PAI_whitepaper_making-ai-inclusive.pdf.

95 Fabio Urbina et al., "Dual Use of Artificial-Intelligence-Powered Drug Discovery," Nature Machine Intelligence 4, no. 3 (March 2022): 189–91, https://doi.org/10.1038/s42256-022-00465-9.

96 Aarathi Krishnan et al., "Decolonial AI Manyfesto," accessed July 24, 2022, https://manyfesto.ai/.

97 Lama Nachman, "Beyond the Automation-Only Approach," in Redesigning AI, Boston Review (MIT Press, 2021), https://bostonreview.net/forum_response/beyond-the-automation-only-approach/.

98 Christina Colclough, "Righting the Wrong: Putting Workers' Data Rights Firmly on the Table," in Digital Work in the Planetary Market, –International Development Research Centre Series (MIT Press, 2022), https://idl-bnc-idrc.dspacedirect.org/bitstream/handle/10625/61034/IDL-61034.pdf.

99 Christina Colclough, "When Algorithms Hire and Fire," International Union Rights 25, no. 3 (2018): 6–7. https://muse.jhu.edu/article/838277/summary.

100 Brishen Rogers, "The Law and Political Economy of Workplace Technological Change," Harvard Civil Rights-Civil Liberties Law Review 55 (2020): 531.

101 Wilneida Negrón, "Little Tech Is Coming for Workers" (Coworker.org, 2021), https://home.coworker.org/wp-content/uploads/2021/11/Little-Tech-Is-Coming-for-Workers.pdf.

102 Jeremias Adams-Prassl, "What If Your Boss Was an Algorithm? Economic Incentives, Legal Challenges, and the Rise of Artificial Intelligence at Work," Comparative Labor Law & Policy Journal 41 (2021 2019): 123.

103 Daniel Zhang et al., “The AI Index 2022 Annual Report” (AI Index Steering Committee, Stanford Institute for Human-Centered AI, Stanford University, March 2022), https://aiindex.stanford.edu/wp-content/uploads/2022/03/2022-AI-Index-Report_Master.pdf.

104 Kofi Yeboah, “Artificial Intelligence in Sub-Saharan Africa: Ensuring Inclusivity.” (Paradigm Initiative, December 2021), https://paradigmhq.org/report/artificial-intelligence-in-sub-saharan-africa-ensuring-inclusivity/.

105 Fekitamoeloa ‘Utoikamanu, “Closing the Technology Gap in Least Developed Countries,” United Nations (United Nations), accessed July 25, 2022, https://www.un.org/en/chronicle/article/closing-technology-gap-least-developed-countries.

106 Annette Bernhardt, Lisa Kresge, and Reem Suleiman, “Data and Algorithms at Work: The Case for Worker” (UC Berkeley Labor Center, November 2021), https://laborcenter.berkeley.edu/wp-content/uploads/2021/11/Data-and-Algorithms-at-Work.pdf.

107 Allison Levitsky, “California Might Require Employers to Disclose Workplace Surveillance,” Protocol, April 21, 2022, https://www.protocol.com/bulletins/ab-1651-california-workplace-surveillance.

108 “The EU Artificial Intelligence Act,” The AI Act, September 7, 2021, https://artificialintelligenceact.eu/.

109 Daron Acemoglu, Andrea Manera, and Pascual Restrepo, “Does the US Tax Code Favor Automation?,” Working Paper, Working Paper Series (National Bureau of Economic Research, April 2020), https://doi.org/10.3386/w27052.

110 Emmanuel Moss et al., “Assembling Accountability: Algorithmic Impact Assessment for the Public Interest” (Data and Society, June 2021), https://datasociety.net/wp-content/uploads/2021/06/Assembling-Accountability.pdf.

111 Kofi Yeboah, “Artificial Intelligence in Sub-Saharan Africa: Ensuring Inclusivity.” (Paradigm Initiative, December 2021), https://paradigmhq.org/report/artificial-intelligence-in-sub-saharan-africa-ensuring-inclusivity/.

112 Daniel Zhang et al., “The AI Index 2022 Annual Report” (AI Index Steering Committee, Stanford Institute for Human-Centered AI, Stanford University, March 2022), https://aiindex.stanford.edu/wp-content/uploads/2022/03/2022-AI-Index-Report_Master.pdf.

113 Business Roundtable, “Statement on the Purpose of a Corporation,” July 2021, https://s3.amazonaws.com/brt.org/BRT-StatementonthePurposeofaCorporationJuly2021.pdf.

114 Larry Fink, “Larry Fink’s Annual 2022 Letter to CEOs,” accessed May 27, 2022, https://www.blackrock.com/corporate/investor-relations/larry-fink-ceo-letter.

115 Katanga Johnson, “U.S. SEC Chair Provides More Detail on New Disclosure Rules, Treasury Market Reform | Reuters,” https://www.reuters.com/business/sustainable-business/sec-considers-disclosure-mandate-range-climate-metrics-2021-06-23/.

116 “Your Guide to Amazon’s 2022 Shareholder Event,” United for Respect, accessed May 27, 2022, https://united4respect.org/amazon-shareholders/.