



# **AI and the Future of Work**

## Sectoral Impacts and Philanthropic Pathways

Amelie De Leon, Jonathan Griffin, James Owen, Juanita Silva, Mariana Valle,  
Pol Villaverde, Rameen Dhindsa, Yuanji Gao, Zi Rui Pang, Zihou Yin

Capstone Project – Columbia SIPA, May 2026



# Table of Contents

*Executive Summary* ..... 5

*Media Narrative Analysis* ..... 7

*Survey Analysis* ..... 10

*Sector Highlights* ..... 15

*Healthcare* ..... 16

*Media* ..... 21

*Entertainment* ..... 28

*Education* ..... 35

*Recommendations* ..... 45

*Conclusion* ..... 49

*References* ..... 53

*Appendix* ..... 57

# Executive Summary

Artificial intelligence (AI) is being integrated into everyday work at an accelerating pace. But its actual impact on workers remains uneven, fragmented, and often invisible in the aggregate data that dominate public conversations. Across industries, there is a growing gap between how widely AI is being deployed and how meaningfully it is transforming work in practice.

Public narratives about AI and work have become increasingly extreme. Executive claims circulate widely, but their supporting evidence is often unclear. In early 2026, Microsoft's AI chief [predicted](#) that AI could displace all white-collar jobs within eighteen months. An [essay](#) by AI entrepreneur and venture capitalist Matt Shumer similarly argued that AI was already automating roles within the industry, [reaching](#) over 80 million views. Dario Amodei, CEO of Anthropic, even [suggested](#) that up to half of entry-level white-collar roles could be displaced within five years. These claims help shape a dominant narrative of rapid and widespread job loss, despite limited empirical evidence to support timelines at that scale.

Against this backdrop, investment in AI continues to grow. In 2024, U.S. private AI investment reached \$109.1 billion, with generative AI alone attracting \$33.9 billion globally (Stanford HAI, 2025). Yet these investments have not translated into meaningful productivity gains. A Federal Reserve Bank of Atlanta survey found that companies reported only a 1.8% productivity increase from AI in 2025, with even smaller gains when measured against revenue data (Federal Reserve Bank of Atlanta, 2026). While AI is being rapidly adopted, its economic impact remains limited and uneven.

Further, many industries have not translated AI investments into operational changes, particularly in highly regulated and unionized workplaces. While 88% of organizations report using AI in some capacity, 81% say it has not had a measurable impact on their bottom line, and only 1% of U.S. C-suite respondents describe their generative AI deployments as fully scaled (McKinsey, 2026). This gap is especially visible at the worker level. AI tools are increasingly present in workplaces but are often introduced without clear guidance or integration into existing workflows. As a result, adoption is unfolding through a mix of top-down deployment and bottom-up, informal use by workers themselves.

Current evidence further suggests AI is reshaping work at the task level rather than eliminating entire jobs. A 2025 study linking occupational AI exposure to U.S. labor data finds that higher exposure is associated with reduced employment and changes in work intensity, particularly among younger<sup>1</sup> and college-educated workers (Dominski and Lee, 2025). Similarly, analysis of U.S. job postings show that demand for routine, administrative roles has declined by 13% since ChatGPT's release, while demand for analytical and creative roles has grown by 20% (Harvard Business Review, 2026). These trends suggest that AI is not uniformly displacing workers, but redistributing demand across tasks and skills.



These dynamics are especially significant for white-collar<sup>2</sup> and pink-collar<sup>3</sup> workers in administrative and entry-level roles<sup>4</sup>. These workers are more exposed to task-level automation and less protected by institutional power, making them vulnerable to how AI is introduced and governed. Research from the UC Berkeley Labor Center (2026) found that AI is more likely to appear in entry-level roles as task intensification, monitoring, or partial automation rather than outright job loss. At the same time, these workers are less likely to have access to training, input into deployment decisions, or protections that shape how these technologies are used, despite being among the most directly affected. However, our research suggests that these workers are not passive. Across sectors, they are already shaping how AI is being used through union negotiations, workplace practices, and individual adaptation.

This project focuses on workers across the healthcare, media, entertainment, and education sectors to better understand how AI is being integrated into everyday work, how its impact is represented in public narratives, and where workers have or lack agency in shaping its use. It asks three core questions:

- 1. To what extent do media narratives about AI and work reflect what is happening in the workplace?**
- 2. How is AI changing day-to-day workplace experience, and what are workers already doing in response?**
- 3. Where are there opportunities for philanthropic and policy interventions that support worker agency?**

Drawing from qualitative research conducted over a six-month period, our research centers the lived realities of frontline workers, with an emphasis on documenting stories of resilience and informal innovation as technologies are introduced into everyday workflows. It consists of a media narrative analysis, semi-structured interviews, and survey data. More than 40 interviews were conducted with frontline workers across the four sectors, providing narrative accounts of how AI tools are being integrated in practice. Lastly, our report incorporates a non-representative survey of over 80 workers, capturing broader patterns in AI exposure, perceptions of its workplace impact, and suggestions for interventions. While not statistically generalizable, the survey surfaces recurring themes and directional insights across a wider set of experiences. In tandem with desk research, our report provides a grounded view of how these changes are unfolding in practice and highlights the role of philanthropy in shaping worker outcomes.

---

<sup>1</sup> We define young workers as any worker below the age of 30.


<sup>2</sup> We define white-collar workers as individuals engaging in “professional, managerial, or administrative work typically performed in an office setting with little to no manual labor” ([Walden University](#), 2025).

<sup>33</sup> We define pink-collar workers as “service-oriented jobs historically associated with women, typically blending aspects of white- and blue-collar job definitions” ([Walden University](#), 2025). Examples include nurses, teachers, and administrative assistants.

<sup>4</sup> We define entry-level workers as individuals with less than a year of working experience who have a basic understanding of the job duties through education or experience ([U.S. Department of Labor](#)).

# Media Narrative Analysis

Our team conducted a systematic media content analysis of 148 articles published between August 2025 and March 2026. Four findings reinforce one another: labor replacement frames dominate coverage across editorially diverse outlets; workers are largely absent from media narratives; narratives reflect a “concerned but measured” tone – acknowledging displacement risk while balancing claims with data or expert debates; and most articles diagnose a crisis without proposing any solution.



## Finding 1: Replacement Dominates the Narrative

Of the 148 articles in our sample, 69 frame AI’s impact on work primarily AI as a direct substitute for human labor. Another 40 use a mixed frame that includes both replacement and augmentation language, meaning that nearly 74% of coverage contains a displacement narrative. Only 28 articles frame AI primarily as an augmentation tool. Think tanks and consultancies (World Economic Forum, McKinsey, Brookings) tend to position AI as a productivity tool with manageable disruption, consistent with their advisory role to the firms deploying it. National newspapers and digital commentary outlets, which depend on audience engagement, lean toward the more dramatic replacement frame. This pattern suggests that the coverage most widely consumed by the public treats displacement as the default story, while the publications offering a more nuanced picture reach a narrower, already policy-engaged audience.


## Finding 2: Workers are missing from the narrative

In 117 of the 148 articles (79%) no affected worker was directly quoted. The outlets with the largest samples illustrate this: Bloomberg includes a worker voice in just 1 of 19 articles, Semafor in 0 of 10, and The Economist in 0 of 7. Even the New York Times includes worker voices in only 3 of 18 articles. The exceptions are Slate, The Atlantic, and Platformer, signaling that smaller, digital-native outlets tend toward more narratives.

When workers do appear in coverage, the category most frequently discussed is “white collar broadly”, a label vague enough to generate anxiety without grounding it in any specific company or role. This category accounts for roughly 43 mentions, followed by “no specific worker group mentioned” at 38, and entry-level workers at around 15. Clerical and administrative workers, who have experienced some of the most measurable employment declines since the launch of ChatGPT in 2022, appear in only a handful of articles, with 12 mentions.

Technology executives are the most cited voice in the coverage, appearing as a source in 55 articles. Economists and academic researchers appear in 53. Affected workers appear in only 25. In 23 articles, no human source is quoted at all. Thus, discussions about the future of work are largely driven by technology companies and researchers, highlighting a gap in worker representation.

#### WORKER GROUPS MENTIONED



#### Finding 3: The tone is measured but the analysis stays shallow

The largest tonal category in our sample is “concerned but measured” (51 articles, or 34%), referring to coverage that sounds reasonable while accepting the replacement frame as given. This narrative tone was then followed by “optimistic” (29 articles, 20%) and “neutral or factual” (26 articles, 18%). Only 20 articles (14%) are coded as “alarmist”, and 8 as “dismissive”.

Labor market evidence does not yet support the scale of disruption implied by the dominant narrative. Many layoffs attributed to AI coincide with familiar patterns: pandemic overhiring corrections, investor pressure to cut costs, and restructuring to redirect spending toward AI investment itself (Marketplace, 2025). Some researchers and journalists have begun describing this gap between claim and evidence as “AI-washing”, companies framing routine layoffs as AI-driven to signal innovation and productivity gains to investors (NYT, 2026; Marketplace, 2026). In several high-profile cases, the

companies announcing AI-driven workforce reductions do not yet have systems capable of replacing the roles being eliminated (Forrester, 2026).

#### **Finding 4: The policy vacuum**

Of the total articles analyzed, 104 (70%) propose no solution at all. Only 39 (26%) do, with 5 offering partial responses. Combined with the replacement-focused framing and the absence of worker voices, this produces a specific effect: the public is left with a narrative of inevitable displacement but lacks a framework to demand action.

The World Economic Forum (WEF) proposes solutions in 10 of 12 articles, and Bloomberg in 9 of 19. But the digital-native outlets almost never do. When policy responses do appear, they tend to be reactive: retraining programs, safety nets, universal basic income, wage insurance for displaced workers. Proactive governance tools, such as the EU Artificial Intelligence Act's risk classifications for workplace AI and pre-deployment transparency requirements, are mostly absent from US media coverage.

The central policy question surfaced by the data is not the threat of labor displacement, but more so, who decides how AI systems enter workplaces and what accountability mechanisms those decisions require. Current media coverage does little to equip policymakers or citizens to ask this question. It centers executive voices, excludes worker perspectives, and offers few regulatory pathways. The narrative around AI and work is moving faster than the evidence, and real people, such as workers, students, and policymakers, are making decisions based on headlines that, according to our interviews, do not reflect what is actually happening on the ground.


# Survey Analysis

To understand how workers themselves are experiencing AI integration, our team deployed a non-representative survey of over 80 workers to identify recurring patterns in AI exposure, workplace impact, and suggested interventions. While not statistically generalizable, our survey surfaces three primary findings:

- ▶ Workers are generally optimistic about AI, but their perceptions of its impact on workload are mixed.
- ▶ Organizations are rapidly adopting AI, while workers are simultaneously resisting, adapting to, and negotiating these changes.
- ▶ Workers express clear demands around legal protection and training, yet employer responses vary in how fully and consistently they address those needs.

**Finding 1: Workers are generally optimistic about AI, but their perceptions of its impact on workload are mixed.**

## PERCEIVED WORKLOAD IMPACT



## SENTIMENT TOWARD AI IN THE WORKPLACE



## HOW MUCH OF YOUR DAILY WORK COULD BE AUTOMATED BY AI?




Our survey respondents currently experience AI as a productivity accelerator rather than a purely disruptive force, though its effects are not uniform. Perceived workload impact is largely neutral-to-positive: 46% of respondents reported that AI and automation somewhat decreased their workload, 18% reported no change, and 17% reported that AI significantly increased their workload. While AI is often associated with efficiency gains, a minority of workers have experience increased burden, likely due to task reconfiguration or learning curves. Sentiment toward AI is also broadly optimistic but tempered. 46% of respondents are “somewhat optimistic” about AI in the workplace, and 24% are “very optimistic”, though a meaningful subset also expressed worries. The data indicates a “cautiously optimistic” stance. While AI is not broadly feared, workers appear to be navigating a transitional phase where productivity gains coexist with uneven workload redistribution. However, this uneven distribution is complicated by workers’ expectations about automation’s reach, with most respondents believing that 40-60% of their daily work could be automated by AI, with the next most common



estimate falling in the 10-40% range. The fact that so many workers anticipate significant automation of their roles, even as their current workloads remain stable or improve, points to a latent tension between present experience and future uncertainty.

**Finding 2: Organizations are rapidly adopting AI, while workers are simultaneously resisting, adapting to, and negotiating these changes.**



High levels of AI deployment across workplaces are evident: 56% of respondents report their employers introducing “many” AI tools, leading to required integration of AI into workflows (54%), and changes in job descriptions due to AI (48%). Despite this, the workers we surveyed are actively contesting AI implementation, with 53% who have challenged how AI will be used in their workplace. Among them, workers have discussed concerns with coworkers, raised concerns with supervisors, and requested training or employer support. As our results demonstrate, workers are actively negotiating AI’s boundaries, primarily through informal peer discussion. Notably, employer responses to the resulting concerns are present, though unevenly implemented: 45% of respondents report that their employers offer legal protections against AI-related layoffs, followed by retraining initiatives (44%), and proactive communication about workplace changes (32%). While these efforts signal growing institutional awareness, their inconsistent reach means many workers are navigating rapid change without adequate organizational support.

**Finding 3: Workers express clear demands around legal protection and training, yet employer responses vary in how fully and consistently they address those needs.**

Even where workers have concerns about AI-driven change, significant barriers prevent them from acting on those concerns formally. 40% believe their employer would not listen, 31% cite lack of time or resources, and 30% fear being perceived as resistant to change, suggesting that structural

constraints inhibit formal engagement with their employer, compounding the organizational gaps identified in Finding 2.


Crucially, what workers say they need most is telling: 60% want legal protections against AI-related layoffs, 53% want retraining, and 51% want clear communication about workplace changes. There is a notable alignment between these priorities and what some employers are beginning to provide, particularly around severance and protections. However, the perceived organizational unresponsiveness limits the effectiveness of these interventions, leaving a meaningful gap between policy intent and worker experience.


These findings paint a picture of a workforce in transition, with workers cautiously open to AI but navigating its rapid integration without consistent institutional support. Workers are actively making sense of a shifting landscape, negotiating change largely through informal channels while facing structural barriers to formal engagement. Our survey ultimately surfaces that organizations are deploying AI faster than trust, communication, and training infrastructure can keep up. Workers perceive this gap acutely, and their demands are, notably, not unreasonable. Legal protections against displacement, investment in reskilling, and transparent communication are foundational workplace expectations. For organizations seeking to realize AI’s productivity potential without eroding workforce trust, the workers most likely to adapt successfully to an AI-integrated workplace are those who feel informed, protected, and genuinely heard.

#### BARRIERS TO FORMAL ENGAGEMENT



#### WHAT WORKERS SAY THEY NEED



The image features a teal background on the left and a close-up of a person's eye on the right. A yellow graphic overlay, consisting of a grid of circles and dots, is positioned over the upper part of the eye. A yellow rectangular box is located at the bottom left of the image, containing the text "Sector Highlights".

## **Sector Highlights**

# Sector Highlights

Finally, to center the lived realities of frontline workers, our team conducted more than 40 interviews with frontline workers across health, media, entertainment, and education. Several cross-sector patterns emerge alongside meaningful differences in exposure, risk, and response.

## Cross-Sector Findings

1. **A near-universal gap exists in employer-provided AI training.** Workers across all sectors reported being largely self-taught, with no role-specific guidance. Workers learned about tools through conferences, peer networks, or experimentation. Where institutions did introduce AI, it was typically deployed without advance consultation.
2. **Existing labor protections do not extend to the workers most exposed to AI-driven change.** Non-union workers, including healthcare staff, freelance journalists, and many in entertainment, report little to no coverage addressing AI-related risks. Even within unionized workplaces, several interviewees were unaware of their own contracts' AI provisions.
3. **A persistent gap appears between formal AI policies and everyday practice.** Many organizations maintain written guidelines that are either unknown to employees or weakly enforced. This disconnect contributes to informal workarounds, including the use of restricted AI tools on personal devices.
4. **Workers identify types of work that remain comparatively resistant to automation,** including clinical judgment in healthcare, editorial decision-making in journalism, and human-centered storytelling in entertainment. The requirement for physical presence also continues to function as a partial buffer against replacement.
5. **What workers worry about reflects their professional identity.** Healthcare workers are concerned about patient safety. Journalists, about credibility, public trust, and compensation. Entertainment workers focused on intellectual property and the extraction of creative labor into training data. Educators worried most about students' capacity for independent thought.

## Divergent Patterns

1. **Perceptions of risk and urgency vary significantly by sector.** Healthcare workers report relatively low personal anxiety, while journalists and entertainment workers express higher levels of concern. These differences appear linked to both exposure to public discourse and the immediacy of observable change.
2. **Experiences with AI tools are mixed and the productivity promise is unevenly realized.** Some workers report efficiency gains while others describe increased workloads due to errors, verification demands, and oversight requirements. Variation appears tied to both task type and the quality of implementation.
3. **Worker responses range from proactive adaptation to disengagement.** A subset of workers is actively reskilling and integrating AI into their workflows, while others acknowledge long-term risks but avoid sustained engagement. What distinguished experimenters from avoiders was access to information and institutional support.

# Healthcare

By Amelie Ortiz De Leon and Juanita Silva

Amid mounting financial and operational pressures, the U.S. healthcare system is turning to AI to sustain efficiency and improve care delivery. Unfolding within a \$4.9 trillion sector, AI adoption is accelerating at more than twice the rate of the broader economy, positioning healthcare as a leading arena for enterprise AI adoption (Menlo Ventures, 2025). This rapid uptake is driven in part by structural pressures: providers continue to face escalating administrative overhead that exacerbates clinician burnout, payers confront rising medical costs, and pharmaceutical firms face extended research timelines and persistently high capital intensity. Given these industry pressures, AI is framed as a solution capable of improving operational efficiency, financial performance, and clinical outcomes.

Beyond these macro-level dynamics, AI is reshaping the everyday practices of healthcare work. For healthcare workers, the technology has the potential to enable “evidence-based and personalized” healthcare (Pailaha, 2023), with use cases involving administrative support, clinical decision support, and interdisciplinary coordination. However, their integration raises important questions about how work is reorganized, how responsibilities shift, and how professional autonomy is maintained. Drawing on qualitative interviews with nine healthcare workers – composed of nurses, unit clerks, a physician assistant, a psychiatric social worker, and a laboratory researcher – we assess AI’s implications for labor organization and care delivery. Below are our main findings:

- ▶ AI adoption is characterized by a disconnect between institutional deployment and frontline use. Several interviewees suggested that familiarity and habitual patterns among workers function as *de facto* barriers to adoption.
- ▶ Labor displacement did not emerge as a major concern in our interviews, with workers distinguishing the salience of labor anxiety by task types (i.e., hands-on patient care vs. administrative work). Worker’s anxieties were centered on care quality, specifically the risk that AI systems would undermine clinical judgment in ways that harm patients, rather than worker displacement.



- ▶ Union engagement is not representative of frontline workers' lived realities, with AI remaining a secondary concern relative to more immediate pressures like staffing shortages and wage stabilization. While workers are well-positioned to engage with AI critically, they lack the training and institutional support to do so on informed terms.

## I. AI Deployment

While our desk research presents AI as a systemic force reshaping healthcare delivery, our interview findings reveal that AI's practical impact remains uneven and fragmented. Across the literature, AI deployment is organized into three domains: administrative support, consisting of workflow optimization tools that reduce clerical burden through automated documentation ([Bracken et al., 2025](#)), EHR assistance ([Microsoft, 2023](#)), predictive staffing ([Palantir, 2025](#)), and algorithmic scheduling used for gig workforce platforms ([Wells and Spilda, 2024](#)); clinical decision support, a layer that enhances diagnostic accuracy through predictive modeling, risk detection, and real-time patient monitoring ([Nanevski et al., 2025](#)); and interdisciplinary coordination, an integration mechanism consolidating fragmented patient data and improving communication across care teams ([Ng et al., 2021](#)). However, our interview findings complicate this framing in three important ways.

First, AI adoption is characterized by a disconnect between institutional deployment and frontline use. Healthcare workers are already engaging with AI tools in informal, individualized ways, most notably through mobile access to applications like ChatGPT. However, institutions are simultaneously restricting consumer tools due to privacy concerns associated with the Health Insurance Portability and Accountability Act ([Zhao et al., 2026](#)), known as HIPAA, and imposing enterprise platforms without sufficient training, integration, or usage protocols. The pattern of AI introduction was identical across every worker we interviewed: tools appeared in hospital systems with little to no advanced notice, announced at most by an email that went largely unread. Among the workers we interviewed, this top-down model had not translated into meaningful adoption. As a physician assistant disclosed:

***“ChatGPT is actually banned in New York-Presbyterian (NYP) locations, so you can't use any of the hospital computers to do anything on AI. But that doesn't stop me. I would go on my phone and find out what I need to.”***

Similarly, an inpatient psychiatric social worker described an AI charting tool appearing in a colleague's system with no email, no training, and no formal policy guidance. As these anecdotes suggest, the presence of AI in healthcare systems does not guarantee its uptake among workers; Instead, blunt restrictions on AI usage encourage workers to bypass institutional constraints.

Second, while our desk research emphasized administrative automation as a major site of transformation, clinicians report little engagement with these systems in their daily work. A clinician noted that AI is absent from their day-to-day practice, remarking that her work was “more hands-

on and involved in patient care.” A registered nurse even explained that when new features were added to familiar systems like Epic, an electronic health record software, uptake remained low:

*“Nurses are creatures of habit, so whatever we're used to, we stick to because we find that easier.”*

Familiarity and habitual patterns among workers function as de facto barriers to adoption, independent of any formal policies.

Third, among the workers we interviewed, AI deployments primarily augmented how clinicians organize and articulate diagnostic information. AI was consistently framed as a cognitive support tool, with interviewees disclosing using enterprise platforms for note taking and transcription purposes. AI was also viewed as most valuable for early-career clinicians, with one respondent describing using AI during neurology consults to identify relevant diagnostic questions, red flags, and structured approaches to patient evaluation:

*“It does really help you when you're fresh out, especially when you're presenting a case to an attending [physician]. You want it to be very straightforward, and AI helps to organize your thoughts in that way.”*

This scaffolding logic extends to administrative communication. A unit clerk described using ChatGPT to professionalize written correspondence:

*“I sometimes struggle to make my emails sound professional. So I'll use something like ChatGPT, write what I want to say, and then rephrase it more formally.”*

Across all roles, the consistent pattern was the use of AI as a support layer for human judgment. Rather than an integrated system transforming healthcare delivery, AI operates as a fragmented set of tools whose most immediate impact is cognitive support and informal, user-driven adoption.

## **II. Labor Anxieties**

Labor displacement did not emerge as a major concern in our interviews, despite the pace of development documented in our desk research. One study found that generative AI can perform on par with non-expert physicians in certain diagnostic tasks, though it continues to underperform in relation to expert clinicians (Takita et al., 2025). Nvidia and Hippocratic AI's healthcare agents were even reported to outperform human nurses on video calls at a fraction of their cost at \$9 per hour, compared to nurses' median hourly pay at roughly \$39 an hour in 2022 (Dumas, 2024). Yet in our conversations with frontline workers, almost none of this landscape was visible to them. Instead, perceptions of AI were structured around task type.

Interviewees consistently distinguished between hands-on patient care and administrative work, using the distinction to assess AI's relevance in their daily responsibilities. Tasks involving direct patient care, which involve physical presence and real-time responsiveness, were widely viewed as resistant to automation. As one nurse explained:

*“To take care of another human, you need to have a human present. To take somebody to the bathroom, to clean them, your physical assessment of the patient has to be with somebody there to see something going wrong.”*

In contrast, administrative functions, like scheduling and documentation, were seen as more susceptible to automation. A unit clerk noted that roles centered on administrative work were more easily automated than those requiring physical movement through the unit, such as nurses and physician assistants.

Further, worker’s dominant anxieties were centered on care quality, specifically the risk that AI systems would undermine clinical judgment and violate HIPPA in ways that harm patients. Our desk research reinforced these claims; A National Nurses United (NNU) survey found that 69% of nurses with employers using algorithmic acuity systems reported assessments failing to match their own clinical evaluations, while roughly 48% of nurses with employers using AI-generated handoffs reported a similar misalignment – with systems omitting patients' educational, psychosocial, and emotional needs ([NNU](#), 2024).

Research on diagnostic tools compound these concerns: ChatGPT Health was also shown to underestimate the severity of medical emergencies in 51.6% of cases and to overestimate urgency in non-critical cases 64.8% of the time ([Ozcan](#), 2026). Our interviewees surfaced the same care-quality concerns in their own clinical experience. A physician assistant recounted the case of a patient who presented with jaw pain, who, based on a known neurological history, she believed an AI system would have classified as trigeminal neuralgia. Her insistence on listening to the patient, who said the pain felt different from usual, led her to order an EKG. It later turned out that the patient was having a heart attack. As she concluded,

***“Medicine isn't always textbook. A patient won't present with a disease like how it's written in a textbook.”***

A nurse similarly raised concerns that an AI system generating patient assignments for charge nurses could not know which nurse was pregnant, had severe anxiety, or, for religious reasons, could not administer blood products. In another case, a psychiatric social worker pointed to a documented case where ChatGPT confirmed a patient's delusion with fatal consequences, arguing that psychiatric care involves a added layer of vulnerability that makes AI errors even more dangerous. Across our interviews, workers articulated their concerns in terms of patient outcomes instead of job security.

### **III. Worker Preparedness**

Given the comparatively slow adoption of AI innovation in healthcare (Apell et al., 2026), frontline workers have a limited awareness of the broader technological landscape encroaching on their professions and responsibilities. While labor unions have positioned themselves as advocates for responsible AI deployment, evidenced by the NNU (NNU, 2025) and NYP strikes (Marcelo, Peltz, and

the AP; 2026) pushing for AI safeguards, their advocacy has barely penetrated interviewees' day-to-day lives. A psychiatric social worker noted that the New York nursing strike included anti-AI provisions, but had not engaged with the substance of those demands. A physician assistant had never heard of any institutional body in her hospital discussing AI. The disconnect between labor advocacy and worker-level knowledge is itself a form of unpreparedness, suggesting that union engagement operates as an exceptional site of awareness that is not representative of frontline workers' lived realities. As one nurse noted:

*“Labor unions are more focused on staffing, wages, and benefits right now. Not AI specifically. During COVID, pay increased significantly due to demand, and it hasn’t fully gone back down. Staffing shortages remain a major issue.”*

This complicates the notion of unions as leading actors on AI, suggesting that even within organized labor, AI remains a secondary concern relative to more immediate material pressures such as staffing shortages and wage stabilization. Thus, as our interviewees noted, healthcare workers are significantly underinformed about the AI tools already in their institutions, let alone what is being developed in the broader sector. The structures that should be informing workers are largely absent, resulting in workers being unable to anticipate, evaluate, or advocate around changes they cannot see. As a researcher at the AI Now Institute noted, the experience of isolation from the employing institution is one of the defining features of algorithmically managed work. Workers want a person they can call when things go wrong, and increasingly, there is no human point of contact. This is the direct result of an information environment that has systematically excluded them.

The workers we interviewed are, in many respects, well-positioned to engage with AI critically: they are experienced, clinically grounded, and already making informal use of AI tools that suggest genuine interest in their potential. What they lack is the training and institutional support to do so on informed terms. Without a significant shift in how healthcare institutions introduce and govern AI tools, the current trajectory points toward a workforce that will be caught off guard by changes well underway.

## **IV. Sector Interventions**

### **1) Establish Protections for Gig Nurses**

Because gig nursing platforms rely on algorithmic management systems to allocate shifts and set wages, formal regulatory classification is a prerequisite for meaningful oversight of AI-driven labor practices. Amid the deregulatory efforts of gig nursing platforms to evade existing labor regulations, policy interventions should register gig nursing platforms as distinct entities (Wells, Pinto, and Spilda; 2026). These responsibilities could include mandating formal registration with the state, quarterly reporting on the number of workers engaged, worker wages, and details on wage and benefit expenditures. New York recently passed a bill (NY S3355) that formally recognizes gig nursing platforms as “entities that must be compliant with healthcare staffing agency rules.” Omidyar can

help advance the regulatory classification of gig nursing platforms through regulatory advocacy, helping establish the preconditions for effective oversight of AI-mediated labor practices.

## 2) Embed Workers Into the Design and Deployment of AI Systems

Worker voices should be embedded in the design and governance of AI systems. While models of AI co-governance primarily center on administrative bodies focused on patient safety<sup>5</sup>, UCLA Health’s AI Council offers a promising model. In 2022, UCLA Health convened a focus group of 20 nurses to evaluate AI tools for nursing workflows. Participants engaged in hands-on validation and provided direct design feedback, ensuring that frontline clinical expertise informed tool development. That feedback is now being translated into iterative deployment: emerging tools like UCLA’s Insight Summaries are currently being piloted with a small group of nurses, who continue to provide input to its developer, Epic. The New School’s Platform Cooperativism Consortium is also exploring an “AI Labor Council for Healthcare,” involving a structured body of clinicians, healthcare workers, and technical staff who evaluate, critique, and shape AI deployment in clinical settings. For a philanthropy like Omidyar looking to build on these emerging models, funding pilot studies and fostering collaboration between tech companies and frontline workers could improve workforce conditions.

---

<sup>5</sup> Examples of these administrative bodies include [Duke Health](#)’s Algorithm-Based Clinical Decision Support (ABCDS) Oversight Committee and [Hackensack Meridian Health](#)’s AI Governance Program.

## Media

By Mariana Valle, Rameen Dhindsa, and Mikayla Pang

The deployment of AI tools into journalism and news media is evolving amid an industry already in structural crisis. Over the past two decades, newsroom employment in the United States has declined, with thousands of local newspapers closing and legacy outlets undergoing repeated rounds of layoffs. The U.S. is losing local newspapers at a rate of over two per week ([Northwestern](#), 2025). AI enters this landscape not as a new disruption, but as the latest shock to an industry that has already been hollowed out by platform-driven revenue loss, declining advertising income, and ownership consolidation. As a veteran journalist and Columbia professor put it:

*“In journalism, we are entering this phase in a very weak position — most organizations have been in crisis for years, the business model has been affected by companies like Google and consumer habits changing. So we are already trying to catch up.”*



Whether AI deployed in news settings is additive, helping journalists produce better work, or extractive, scraping journalism content to feed models that then compete with it, remains an open question. AI companies have used news as training data, and some, including OpenAI, have simultaneously entered [licensing deals](#) and [funding relationships](#) with outlets, creating a dependency in which the same companies undermining journalism's business model are also positioning themselves as its benefactors, repeating the same pattern of the social media platforms.

AI tools are being adopted across newsrooms at an uneven pace. Major news organizations, including the New York Times, the Associated Press, the Washington Post have invested in initiatives to explore how AI can be integrated alongside human journalists ([IBM](#), 2025). Yet when media managers are asked to characterize the impact of current AI initiatives, only 13% describe them as transformational; 44% say promising, and 42% say limited ([Reuters Institute](#), 2026). AI-related job cuts in journalism also remain marginal for now: around two-thirds of media managers surveyed by the Reuters Institute said they have not cut any jobs due to AI ([Reuters Institute](#), 2026). This may, however, be starting to change: in April 2026, the Associated Press offered buyouts to more than 120 staff members as the company signaled a pivot toward AI and video ([Fortune](#), 2026).

Much of the discourse around AI in media centers on journalists at prestigious, legacy outlets, who have the institutional power and public visibility to shape the conversation. This framing risks marginalizing freelance journalists, local news reporters, translators, copy editors, fact-checkers, and production staff, whose work sustains the broader information ecosystem but who are absent from evaluative research and AI tools design. In 2024, 34% of U.S.-based journalists worked as freelancers or are self-employed ([Pew Research](#), 2024), and another research from 2022 reveals that only 16% of American journalists at news outlets are unionized ([Pew Research](#), 2022).

## I. AI Deployment

AI applications in journalism are mostly concentrated in three domains: content production, news gathering and investigation, and curation, distribution and personalization. News organizations already used machine learning for tasks like social media monitoring and data management ([IBM](#), 2025), but Generative AI's new capacity to produce written content, audio, video, and images raises fundamentally different questions about the role of human journalists.

- **Content Production:** This is where AI's role in journalism becomes most contested. AI is being mostly used for summarization, translation, simplification, rewriting in different styles, and extracting copy for social media, newsletters, and scripts ([Reuters Institute](#))., but several outlets have gone further: Fortune and Business Insider have explored using AI to write full article drafts, notifying readers of their intent to use AI for this purpose ([NYT](#), 2025); McClatchy, parent company of the Miami Herald and others, deployed a “content-scaling” system designed to

produce short- and long-form summaries of reporters' stories, create audience-specific versions, and generate scripts for short-form video ([The Wrap](#), 2026).

- ▶ **News Gathering and Investigation:** AI tools also support the investigative and analytical dimensions of journalism. The Associated Press used AI tools to sort through tens of thousands of pages of declassified documents related to the assassinations of JFK, MLK, and RFK, making documents searchable, summarizing them, and identifying newly unredacted sections; work that the AP's director of AI products estimated saved reporters "days' worth of work" ([NYT](#), 2025). Analysts predict that AI can alleviate "tedious" operational tasks: tagging, categorizing, adding metadata, headline and SEO suggestions, copyediting, organizing research, processing permissions, and moderating comments ([Reuters Institute](#)).
- ▶ **Curation, distribution and personalization:** News organizations have used AI to analyze audience engagement and track trending topics for years ([IBM](#), 2025). More recently, AI has been applied to personalized newsletter delivery, dynamic paywall optimization, A/B testing of headlines, push notification targeting, and automated comment moderation. These tools operate at the distribution layer, shaping how audiences encounter journalism rather than how journalists produce it. None of our interviewees described awareness of or interaction with these systems, which indicates that the AI tools closest to the business and product side of news organizations are the ones least visible and most distant to the workers producing the journalism, and usually more exposed to the leadership making decisions.

The literature frames AI as a tool that can transform newsroom workflows end-to-end and major outlets have invested in AI teams, AI-specific editorial roles, and partnerships with technology companies. Drawing on nine qualitative interviews with media workers – including journalists (reporters, freelancers, leaders), professors, media policy experts and researchers – our findings reveal that AI's practical impact on journalism remains limited and fragmented.

First, the journalists described AI use that is overwhelmingly informal, self-taught, and limited to a narrow band of pre-writing and research tasks. A freelance journalist in Los Angeles described her AI use as confined to transcription (using Apple's Voice Memos app) and document analysis (using Google Pinpoint for Journalists to search lengthy city council agendas). A freelance foreign correspondent covering conflict in Ukraine and Gaza described using Otter AI to transcribe interviews and prompting an LLM to extract quotes and identify contrasts across sources, work that "would have originally taken me 5+ hours".

Second, we found that the promise of AI improving newsroom productivity is partially true. For some workers, using AI summarization requires so much additional fact-checking that it consumes more time than doing the work manually. A journalist we spoke to described a systemic problem:

*"There is a push from higher-ups to use AI to increase productivity. However, there is a mismatch in the perception of what AI can do versus what it can actually do. AI does not actually always improve productivity — AI introduces more error, which creates more work for editors."*

Interviews with journalists consistently revealed that while AI tools are valuable for managing repetitive administrative tasks, conducting initial research, and processing large datasets, the inherent risk of algorithmic hallucinations and factual errors necessitates that human journalists remain strictly responsible for vetting all final content.

Media narratives and research mostly reflect AI vendors' and newsroom executives' claims about AI's capabilities, while interviews show actual user experiences. The gap matters because productivity gains justify AI adoption. If these gains are illusory or uneven, saving time for reporters but causing errors for editors, then the productivity rationale weakens, reducing costs at the expense of editorial quality.

A researcher at Harvard we interviewed provided the structural explanation for this mismatch. Much of what companies frame as AI deployment is simply “an enterprise version of ChatGPT”. The result is that the AI tools being pushed onto journalists were not designed for journalism. “They are general tech like AI is being retrofitted to the workplace as a top-down approach, rather than purpose-built tools designed around actual workflows.”

Third, we found a distinction between administrative and creative AI use. Workers across our sample were comfortable with AI performing what one interviewee called “low-effort, low-impact” tasks: transcription, document search, formatting, and basic research synthesis. But they drew a sharp line at anything that would, in a journalist’s words, “stand in as me.” The foreign correspondent was explicit: “AI can’t build trust with people, speak their language, eat with them, spend time with their families, see how war and atrocities have actually affected them.” A reporter extended this logic to the production workflow: AI could handle idea generation, research, and editing, but the distinctly human act of reporting, going to the field, exercising judgment, building relationships, remained non-substitutable.

Every journalist in our sample drew the same line between acceptable and unacceptable AI use, and what AI can do and what only a human journalist can do. However, as the technology evolves and models become more capable, the line moves faster, and some outlets are already crossing what the journalists considered safe. McClatchy’s content-scaling agents producing articles under reporters’ bylines, Politico’s AI-generated live political coverage, Axios’s automated local news roundups are some examples we documented.

## **II. Labor Anxieties**

Workers articulate concerns based on what they can directly observe, not on the full trajectory of AI deployment in the media industry. Major concerns are around journalism quality and displacement, but they anchored their anxieties in specific, role-dependent fears rather than a general sense of threat.

*“My mantra is either we train AI into our system to help us produce better journalism, or AI is going to take our jobs.”*

Workers consistently expressed confidence that on-the-ground reporting, source-building, and editorial judgment were non-substitutable, but that AI will eliminate entry and mid-level roles, the training pipeline through which junior journalists develop those skills. “When it comes to entry-level roles, we’re going to lose a lot, same with mid-level editing jobs.” The foreign correspondent identified the same vulnerable population: “Probably fact-checkers, support staff, entry-level journalists without a following.” The journalist covering the media industry confirmed this from the employer side: “There is no talent pipeline because entry-level roles are hiring less and less.”

A second cluster of anxiety centers on credibility and public trust. The Daily News reporter raised objectivity and representation: if AI tools are trained on biased data, they may reproduce political leanings and further undermine public trust in an industry already facing a credibility crisis. The McClatchy case illustrates the credibility dimension: when AI produced articles under reporters’ bylines without their knowledge, the reporters’ concern was not primarily about their own jobs but about the absence of consent, their credibility and the betrayal of public trust. As investigative reporter Ariane Lange told The Wrap: the tool was presented as an experiment, but “the imperilled guinea pig is our credibility” ([The Wrap](#), 2026). At the New York Times, editorial union leaders told management that its AI standards were “too vague and inadequate,” creating editorial problems and trust issues ([Axios](#), 2026).

The Daily News reporter expressed concern that AI adoption could disproportionately harm journalists from marginalized backgrounds, such as trans and LGBTQ journalists, who already face discrimination. AI automates entry-level roles, often filled by early-career journalists from underrepresented groups relying on these positions to enter the profession. Automation thus impacts diversity and inclusion efforts. Additionally, if AI tools are trained on biased data, it further narrows the variety of stories and perspectives. As newsrooms shrink, those with less institutional power, many from marginalized communities, are the first to lose jobs. This critical dimension of AI and labor is often overlooked, with focus mainly on overall job numbers and productivity.

Another concern from many workers is of the industry falling behind. While they acknowledge the risks of using AI, they are convicted that resistance to the technology is itself a risk.

*“When I was working as a reporter, we were driving around with a flip phone and every car had a map with the five boroughs. That’s like me saying I want to use paper maps instead of Google Maps.”*

The foreign correspondent made a parallel point: “Lots of people in the industry want to act like they don’t use it or will never use it, but that’s just a disservice to themselves.”

### **III. Worker Preparedness**

We interviewed two journalists specializing in the development of AI tools for their respective organizations to understand how these technologies are being integrated into journalistic workflows. Both interviewees emphasized that the most effective AI tools are co-designed with direct input from

the journalists who will eventually utilize them. By involving the workforce in the creation and development process, organizations can more effectively prepare staff for technological transitions. Furthermore, these journalists receive formal training that establishes clear guidelines, defines permissible use cases, and sets rigorous limitations and expectations for AI performance.

For instance, one of our interviewees stated that:

*“The training has been designed by journalists for journalists and is mandatory.”*

This emphasizes that involving the end-users throughout the entire development process is vital. Journalists themselves are best positioned to identify which operational gaps AI can effectively bridge and which core tasks must remain exclusively human.

However, outside of major national outlets, AI adoption is driven almost entirely by individual initiative rather than institutional strategy. The freelance LA journalist learned about Google Pinpoint at a journalism conference; the foreign correspondent taught himself to use Otter AI and Claude; the Daily News reporter had no organizational-level conversation about AI at all.

Purpose-built tools, like Google Pinpoint, which helps with efficiently search public records, are delivering real value to the reporters who find them. The problem is not that useful AI tools for journalism don't exist, but our findings suggest that newsrooms may be investing in the wrong tools through the wrong process. They do not talk about AI, or adopt a top-down deployment of general-purpose technology without worker input, rather than identifying, with the journalists, the specific tools that match the actual work. While some rules do exist, they are unevenly applied and poorly communicated. A freelance journalist noted that the newsroom she works for has a strict no-AI policy that prohibits freelancers from using AI note-taking tools in meetings, yet staff writers have access to an internal AI chatbot that pulls from the company's style guide.

The information asymmetry extends beyond AI tools to the commercial deals that govern how journalists' work is used. When newsroom management strikes a licensing agreement with an AI company, the terms, what content is included, how it will be used, and compensation, are typically unknown to the journalists who produced the work. As one expert with extensive union and policy experience declared: 'The management knows things that the journalists do not.' This is a different order of information gap from not knowing which tools will be deployed. It encompasses the commercial use of workers' output without their knowledge.

Unions in media and entertainment sectors are the most active in bargaining AI deployment. The NewsGuild-CWA has secured AI protective language in more than 30 newsroom contracts and won a landmark arbitration ruling at Politico establishing that management must bargain before deploying AI tools ([Wright & Grohmann, 2026](#)). The Daily News reporter described her union contract as providing meaningful security: “Certain persons can never be replaced by AI.” She saw the union as a “trusted party to communicate with the corporation”, a channel for information and advocacy that no other institutional structure provides. But this protection is structurally limited. Freelancers,

who constitute a significant and growing share of the journalism workforce, as noted before, have no equivalent, they operate in a gray zone where organizational AI policies may or may not apply to them, and where they have no seat at any bargaining table.

#### IV. Sector Interventions

According to a media policy expert we interviewed, the challenge of protecting journalism workers in the age of AI is complicated by a three-way misalignment of incentives that distinguishes media from other sectors. Publishers and AI companies are fighting over who captures the economic value of journalism through licensing deals, and lobbying, but this fight is between institutions, not on behalf of workers. When a publisher secures a licensing deal, the revenue flows to the organization, not to the reporters whose work was included in the training data<sup>6</sup>, and they don't know the terms of the agreement; meanwhile, management deploys AI internally to cut costs, and workers push back through unions and grievances. Workers are caught between both conflicts without structural leverage in either.

The policy space reflects this misalignment. Publishers are well-funded and actively shaping legislation to protect their revenue through licensing compensation, tax credits, and subsidies. Technology companies are well-funded and actively lobbying for less regulation. Both parties have dedicated policy teams, lobbyists, and institutional infrastructure. Workers have the unions, and the unions are doing consequential work, but they can only bargain for their members, who represent 16% of the journalism workforce. Beyond the guilds, as the expert put it, 'there are no incentives for policy development that's worker-centered.' He mentioned that no organization is funded to work on the structural labor conditions that affect all journalists: union and non-union, staff and freelance, national and local.

Philanthropy has stepped in to address journalism's economic crisis, funding newsrooms, subsidizing reporting, and supporting AI adoption through programs like the IPI Global AI Accelerator, which provides small and independent news outlets with funding, training, and advisory support to develop AI tools and workflows. Although they are oriented toward newsrooms as institutions, not to the labor conditions, they add a meaningful contribution: they empower journalists to shape how AI is integrated into their work rather than having it imposed by management or technology companies. The worker-centered dimension remains largely unfunded.

The interventions below are designed to build the policy infrastructure and worker power needed to change the conditions that we found:

- 1) **Fund worker-centered policy development:** The policy space is currently dominated by publishers seeking revenue protection and tech companies seeking less regulation. A dedicated policy team, housed at an existing media/labor policy center, would draft model

---

<sup>6</sup> In June 2024, Le Monde agreed to [redistribute](#) 25% of all AI licensing revenue directly to its staff journalists as a yearly bonus – a landmark move.



workers' board legislation at the state level, and produce analysis on how to articulate sectoral bargaining, and build the legislative relationships needed.

- 2) **Build shared AI contract infrastructure:** The NewsGuild has secured AI protective language in more than 30 contracts, but the clauses are scattered across individual union websites and difficult to compare. A searchable, structured database would let unions benchmark future negotiations and give non-union and freelance journalists visibility into what protections are possible, even without formal bargaining power. CLCF at the University of Toronto already maintains a worker mobilization tracker, but no information on the contract terms.
- 3) **Create a sustained AI information channel for journalists:** Create a knowledge hub with newsletter, webinars, and a practical guidance on AI tools, rights, and protections, specifically targeting freelancers and journalists at small/local outlets who have the least access to institutional resources.

## Entertainment

By Jonathan Griffin, James Owen, and Zihou (Silver) Yin

The entertainment industry is one of the sectors most visibly shaped by rapid advances in artificial intelligence (AI) and automation. Existing literature often presents a precarious outlook, warning that entertainment workers, from sound editors to actors, may face challenges related to job security, compensation, intellectual property, and transparency. While these concerns are real, the story is more complex than simple technological displacement.

The main impacts of AI in the entertainment industry center on four areas. First, task automation is reshaping creative workflows, particularly in writing, editing, and visual effects, raising concerns about deskilling rather than outright job loss (Demirer et al., 2026). Second, consent issues have intensified, especially regarding the use of actors' digital likenesses without clear consent or compensation (SAG-AFTRA 2023 contract summary). Third, AI complicates traditional compensation models, particularly residuals tied to human-created content. Finally, there is a growing lack of transparency, as workers are often unaware of how AI tools are integrated into production processes. Importantly, interviews with industry professionals suggest that the immediate concern is less about mass unemployment and more about declining bargaining power and weakening labor standards.

Although consent and compensation are often discussed together, they are distinct: consent governs whether a worker's labor or likeness can be used by AI at all, while compensation governs how workers are paid when it is.

These anxieties are not without precedent. The entertainment industry has repeatedly adapted to technological disruption. The transition to sound in film during the late 1920s displaced some silent film actors while creating new roles (Crafton, *The Talkies*, 1997). The rise of television expanded the industry despite initial fears of decline. More recently, digital editing and CGI required workers to upskill, while streaming platforms like Netflix disrupted compensation structures, particularly residual payments (McKinsey Global Institute streaming report). Across these shifts, technology has tended to reorganize labor rather than eliminate it, though often at the cost of worker leverage (Brynjolfsson & McAfee, 2014).

Emerging worker-led coalitions and informal networks further illustrate how responses to AI are evolving. These groups focus on transparency, fair crediting, and ethical standards, often operating alongside traditional unions. Overall, while media narratives often frame AI as an existential threat to creative labor, the reality is more nuanced (Brynjolfsson & McAfee, 2014; Acemoglu & Restrepo, 2018). The central issue is not simply whether AI will replace workers, but how it will reshape power dynamics within the industry. As past technological disruptions suggest, the future of work in entertainment will depend less on the technology itself and more on how workers, institutions, and policymakers negotiate its use.

Drawing on qualitative interviews with entertainment workers spanning music production, film and television sound, post-production editing, and acting, we assess AI's implications for labor organization and creative production. Below are our main findings:

- ▶ AI adoption in entertainment is uneven and largely informal, concentrated in pre- and post-production rather than on-set creative work, and driven more by individual initiative than institutional strategy.
- ▶ Labor displacement anxieties are real but nuanced: workers are less concerned with outright job elimination than with the erosion of entry-level pipelines, the devaluation of specialized skills, and the quiet absorption of human creative labor into AI training datasets without consent or compensation.
- ▶ Workers are significantly underprepared for the pace and scope of AI's integration. Unions have made important gains, but their protections remain structurally limited and have not yet translated into meaningful awareness or security for non-unionized and freelance workers.

## **I. AI Deployment**

Across the literature, AI tools in entertainment are organized primarily around three domains: content generation, which includes script drafting, image synthesis, and automated voiceover; post-production augmentation, covering editing assistance, sound mixing, and visual effects workflows;

and distribution and discovery, encompassing algorithmic curation, recommendation systems, and audience targeting (McKinsey Global Institute, 2023; World Economic Forum, 2025). However, as in healthcare and media, our interview findings reveal a substantial gap between the capabilities described in industry reporting and the tools workers actually encounter in their daily practice.

First, AI adoption is highly uneven across production roles. Workers in on-set and performance-facing roles report minimal direct engagement with AI tools. An actor we interviewed described showing up to set without having encountered AI in any form:

*"I haven't seen it incorporated. At the moment... if I were to show up on set, I haven't encountered any use of AI."*

The same interviewee noted that AI in their sector tends to appear either upstream, in development and pre-production, or downstream, in post-production and distribution — rarely in the visible, collaborative space of active production. For workers whose labor is centered on performance, presence, and relationship-building, AI remains largely abstract.

Second, where AI is present, its impact is concentrated in specific, task-level functions rather than wholesale workflow transformation. In post-production sound, the most clearly disrupted area identified by our interviewees, AI tools are being used to complete narrow, technically defined tasks — particularly those involving language processing and dialogue (Demirer et al., 2026). One industry committee member who formally evaluates AI applications for a union described the pattern precisely:"

*"There are tools that complete very specific tasks and will do those very well, and for jobs that handle those tasks right now, that creates a certain level of occupational threat, but industry as a whole — not a huge upheaval."*

At the same time, they pushed back on catastrophist framings, noting that broader job losses in their guild have other causes. Crucially, they situated current AI shifts within a longer history of crew reduction:

*"There were 50 people working on sound [on a major film in 1995]. Today that type of project would have maybe six people."*

Third, AI's reach in entertainment is also felt acutely in the music industry, where the displacement of session and demo musicians represents a concrete, ongoing economic harm — one that does not register as a headline layoff but steadily erodes livelihoods (Acemoglu & Restrepo, 2018). A music producer described the structural shift:

*"People who used to do demos... that's declining significantly. Not only the producer who would have put the session together, but also every musician in that room that would have helped*

*realize that idea. They are no longer getting as many calls for that kind of work as they would have even five years ago, let alone 10 or 15 or 20 years ago."*

The same interviewee noted that AI-generated music is already displacing human composers in the sync licensing market — the space where music is created for commercials, film, and television. Streaming platforms further compound this dynamic. One interviewee described discovering AI-generated artists and tracks appearing alongside human musicians on algorithmic radio playlists:

*"I would go onto my Spotify page... it was me and a couple human artists and a lot of AI generated bands and artists... a lot of the AI stuff, they kind of have a watery quality to the audio, because it's not real."*

The presence of AI-generated content into distribution infrastructure undermines the discoverability and income of human artists, and represents a form of displacement that is diffuse, hard to quantify, and largely invisible in the aggregate labor statistics that dominate public debate (Brynjolfsson et al., 2025). Finally, post-production cost reduction is being actively pursued through AI tools in film and television, particularly for reshoots. One industry professional described a growing practice:

*"There's been a big push with a few companies specifically to cut down on the need for doing reshoots... instead of having to pay \$20,000 for a couple days' reshoots, you pay half to a quarter of that using a specialized tool to make the tweaks that you need."*

AI as a cost-compression tool applied to specific production problems is a consistent use case across sectors and may cumulatively reduce the volume of paid work available to human workers, even where no single role is formally eliminated (McKinsey Global Institute, 2023).

## **II. Labor Anxieties**

Entertainment workers' concerns about AI are anchored in specific, role-dependent observations rather than generalized alarm. Across our interviews, a recurring pattern emerged: workers drew sharp distinctions between the tasks AI can competently perform and the inherently human dimensions of creative work. They simultaneously acknowledge that the line between the two is moving faster than their institutions can track. The clearest and most consistent anxiety centered not on the immediate elimination of senior creative roles, but on the destruction of the entry-level pipeline through which workers have historically developed those skills (Brynjolfsson et al., 2025). Multiple interviewees identified junior, support, and freelance positions as the most immediately exposed. A performer and voice actor observed: *"I do think there will be fewer opportunities in those kinds of roles."*

The devaluation of specialized craft is a second distinct anxiety. Entertainment workers across our sample expressed concern not just about job loss, but about the intellectual capital embedded in their labor being extracted and monetized without consent or compensation. Unlike a layoff, this form of displacement is largely invisible in aggregate labor statistics. The worker's labor persists in the model,

but they receive no credit and no residual benefit. One interviewee described the concern in specific terms:

*"There is a company that could take sound mixes across hundreds of films... take all that work, that intellectual capital, and turn it into a model and then profit on that model in perpetuity."*

A third cluster of anxieties centers on likeness, voice, and consent. For performers, the threat is not only economic but existential in a distinct way: their identity, their face, their voice, their physical presence, can be replicated, reused, and deployed without their knowledge. One actor commented:

*"Studios using your likeness or voice without consent, even after you've passed away... contracts from 20 years ago didn't account for this technology, so there's no protection in place for older work."*

The concern extends to the posthumous use of a performer's likeness, which raises ethical questions that existing labor law has not yet addressed. This dimension of AI's reach — the potential to sever the link between a worker's identity and their compensation entirely — may mark a qualitative shift from prior technological disruptions (Acemoglu & Restrepo, 2018). Despite these anxieties, resistance is not workers' dominant stance. Several interviewees pushed back against reflexive rejection of AI, framing adaptation as a professional necessity:

*"Lots of people in the industry want to act like they don't use it or will never use it, but that's just a disservice to themselves."*

The tension between critique and adaptation runs through all of our interviews, and reflects the broader uncertainty of a workforce navigating a technological transition without adequate institutional guidance (Brynjolfsson & McAfee, 2014).

### **III. Worker Preparedness**

Entertainment workers are navigating AI's integration largely without formal institutional support. Across our interviews, no interviewee described receiving structured training, a clear organizational AI policy, or proactive guidance from an employer. Where AI tools were being used, adoption was almost entirely self-directed: workers discovered tools through peer networks, conferences, and individual experimentation, and made decisions about use and limits on their own. This pattern mirrors findings from across our sectors and is consistent with research on top-down AI deployment without worker input (Pailaha, 2023).

This pattern holds even within unionized workplaces, where AI governance has advanced furthest. The unions most active in this space — SAG-AFTRA, the WGA, and IATSE — have secured important contract language around AI consent, likeness protections, and the requirement to bargain before deploying new tools. One interviewee with a union contract described it as providing genuine security that certain persons can never be replaced by AI. The most significant preparedness gap is structural rather than individual. Workers in music, freelance production, and non-union performance roles

have no equivalent to the contract protections that unionized workers have fought for. Freelancers operate in a gray zone where organizational AI policies may not apply to them and where they have no seat at any bargaining table. This structural exclusion mirrors the freelance divide documented in the media sector (Pew Research, 2024).

Within unions, there is a further gap between the protections that have been won and workers' awareness of those protections. Just as frontline healthcare workers were largely unaware of the AI provisions in their union contracts (NNU, 2024), entertainment workers in our sample showed uneven knowledge of what their contracts actually covered. An actor noted that specific AI provisions were still being negotiated, and that the process was opaque:

*"I haven't personally seen anything specific about AI yet [in contracts]. I think it's because negotiations are still ongoing."*

The committee-based approach adopted by some unions — where dedicated working groups formally evaluate AI tools, stress-test their capabilities, and make recommendations to governing boards — represents a meaningful model for institutionalizing worker voice in AI governance. One interviewee described their committee's work:

*"What we do as a committee is we analyze these applications and break them apart and stress test them, figure out what they can do, what they can't do, and then make recommendations to our board of directors."*

Critically, this committee had already identified the most exposed roles in their jurisdiction and was proactively preparing a response before displacement became acute. But this level of institutional investment is exceptional, not representative of how most entertainment workers encounter AI.

## **IV. Sector Interventions**

Upon distilling the wide array of responses from our interviews, three themes become clear: research into model attribution and compensation mechanisms, streaming transparency, and cross-sector coalitions on training data.

### **1. Worker Compensation**

The use of human creative labor to train AI models without consent or compensation is one of the most structurally significant harms surfaced in our interviews, and one that remains largely unaddressed in current policy frameworks. A sound mixer's work across hundreds of films, a session musician's catalog, a voice actor's recordings — this labor can be absorbed into a model that then competes with the worker for future jobs. The output persists indefinitely, but the worker receives no credit, no royalty, and no residual benefit. Philanthropic funding can help build the infrastructure to change this: research into model attribution and viable compensation mechanisms, support for coalitions linking entertainment workers with visual artists and writers facing parallel challenges, and



the development of model contract language that establishes consent and compensation as defaults rather than exceptions. The Le Monde precedent cited in the Media section, where 25% of AI licensing revenue is redistributed to staff journalists, offers one transferable model, though residual structures and union jurisdictions in entertainment will complicate direct adaptation. Progress here will be incremental, but the precedent that training data is freely available for extraction is hardening quickly, and the window for establishing a different default may be narrowing.

## **2. Establish a Sync and Streaming Transparency Standard**

One of our interviewees floated the idea of a mandatory "Made with the help of AI" stamp or label that would apply across entertainment industries at the state level. The thinking is pretty straightforward: audiences, and arguably workers too, deserve to know when what they're watching, listening to, or reading was made with AI involvement. This came up specifically in the context of music, where AI-generated tracks are already showing up on streaming playlists next to human artists without any disclosure, but the same logic applies to film, television, and voiceover work. We heard versions of this concern across other sectors as well, which suggests labeling is one of the few interventions with broad cross-industry support.

A state-level approach makes sense given how slow federal action has been. California or New York could move first, given that most entertainment production runs through them, and other states would likely follow. The harder question is what counts as "made with AI," a fully AI-generated song is obvious, but what about a track where AI was used for mastering, or a film where AI handled some reshoot tweaks? Any policy would need to figure out where the line is, ideally with input from the workers actually doing this work day to day.

## **3. Build Cross-Sector Coalitions on Training Data and Intellectual Capital<sup>7</sup>**

A recurring concern across our interviews was that the intellectual capital embedded in workers' labor, sound mixes, performances, scripts, compositions, is being absorbed into AI training datasets without consent, credit, or compensation. This is not a problem unique to entertainment. Journalists, visual artists, voice actors, and authors are facing structurally identical challenges, and current advocacy efforts are fragmented across each of these groups. Philanthropic funding should support coalition-building that brings these constituencies together to develop shared legal and policy strategies on training-data consent, compensation, and the right to opt out. A coordinated push carries more weight than parallel efforts moving separately, and the legal questions, what counts as authorized use, how compensation should be structured, what enforcement looks like, are largely the same across sectors.

---

<sup>7</sup> More broadly, cross-sector coalitions are emerging as a strategy to build collective leverage across industries. For example, National Nurses United organized [solidarity rallies](#) in support of striking writers and actors, highlighting how labor groups are aligning across sectors to contest the use of AI and advocate for worker protections.

# Education

The education space is often regarded as one of the areas where the stakes for societal changes driven by technology are highest, given how human capital ultimately underpins what is possible across the economy. On the one hand, it is socially expected that the education system will prepare the next generation of the workforce to master new technological tools enabling the work of the future. On the other hand, the way in which education is delivered in itself is also subject to the impacts of AI, with an often significant discrepancy between the techno-optimistic expectations forwarded by edtech firms and the reality in the classroom. This phenomenon broadly labelled as “the educationalization of social problems” ([Smeyers & Depaepe, 2008](#)) transfers the responsibility to deal with nearly everything to schools.

However, rather than increasing the burden on the school system itself, the question of AI in education should be examined *vis à vis* larger social patterns. The classroom setting is deeply interconnected with broader social inequality patterns and the ecological factors shaping classroom policy, teacher skills, and the perception from different actors in the sector. In this line, systemic patterns in school funding and the curation and distribution of AI implementation guidelines will be covered below.

The public narrative around AI in education has been dominated by edtech firms and institutional policymakers promising personalized learning, efficiency gains, and workforce preparedness. Microsoft claims that 86% of education organizations globally now use generative AI, with AI fluency framed as an imperative for both educators and students ([Microsoft, 2026](#)). Survey data suggests rapid uptake: nearly 60% of U.S. public school teachers reported using AI during the 2024–25 school year, with regular users claiming time savings of up to six hours weekly ([Gallup & Walton Family Foundation, 2025](#)). In higher education, 45% of faculty now use AI tools, up sharply from 2023 ([Cengage Group, 2024](#)).

In the following pages, we will dive into both questions while focusing on the perspectives of how education professionals are being affected by AI. Drawing on qualitative interviews with seven education workers and stakeholders, spanning K-12 classroom teachers, vocational instructors, edtech developers, and philanthropic program officers, we assess AI’s implications for instructional labor and the human capital pipeline. Below are our main findings:

- ▶ AI adoption is bifurcated and largely self-directed, with educators improvising use in the absence of clear institutional policy, training, or pedagogical guidance. The tools most commonly adopted are lightweight administrative aids rather than integrated instructional systems.
- ▶ Labor anxieties center on educational quality and human connection rather than job displacement. Teachers consistently distinguish between tasks AI can perform and the inherently relational and developmental dimensions of teaching that resist automation, while expressing acute concern that AI is eroding students' capacity for independent thought.
- ▶ Worker preparedness is structurally unequal and ecologically fragmented. The variance in readiness between districts and individual educators reflects systemic funding disparities, decentralized governance, and the absence of coherent frameworks for professional development.

## I. AI deployment

Across the literature, AI applications in education are organized into three domains: administrative and instructional support, encompassing lesson planning, rubric generation, individualized education program (IEP) documentation, and parent communication; personalized learning and tutoring platforms, including differentiated instruction systems, automated feedback tools, and adaptive learning pathways; and student-facing productivity and academic integrity tools, ranging from writing assistants and code generators to plagiarism detection and process-documentation requirements. However, our interview findings complicate this framing in three important ways.

First, AI adoption is characterized by a disconnect between institutional deployment and frontline use. While education organizations may report high rates of “AI use,” what this means in practice varies wildly and often occurs without formal policy, training, or integration into existing workflows. As of late 2023, UNESCO reported that fewer than ten percent of schools and universities globally had established formal policies on generative AI by the end of 2023 ([UNESCO, 2023](#)). Our interviews reinforced this gap. A high school computer science teacher, described his school's policy requiring students to declare AI use as functionally unenforceable. As he explained,

*“It’s actually hard to enforce. What I do now is require handwritten code for in-class assignments—no computers, just like a math exam. For homework, AI is allowed, but they must submit a ‘process document’—screenshots of their AI conversations, plus a reflection explaining what suggestions they accepted or rejected.”*

A third-grade teacher in Texas noted that her principal used AI to draft internal communications, a practice that eroded trust rather than modeling constructive integration:

*“When your leadership is relying on a robot [...] if you couldn’t spend 3 minutes to write a good morning message about the events that are coming up, then we have a problem.”*

The pattern closely parallels findings from healthcare and media: tools appear in institutional settings with little advance notice, no training, and weak enforcement, producing a vacuum in which educators either reject AI outright or improvise use on personal initiative.

Second, where AI is present, its impact is concentrated in pre-instructional administrative and surveillance tasks rather than in core pedagogical delivery. Some teacher’s grading workflow exemplifies this displacement: their evaluation has shifted from assessing student logic and syntax to detecting whether code was AI-generated, a task that consumes additional time rather than saving it. As he disclosed,

*“Before, I only looked at the final code when grading. Now I have to review chat logs. And honestly, some students fabricate their conversation records to get around the requirement, or deliberately ask AI stupid questions to make it look like they were ‘guiding’ the AI. Identifying these takes time too.”*

Similarly, a learning innovation researcher at the University of Pennsylvania, noted that while teachers use AI for rubrics, assessments, and lesson planning, she always emphasizes that teachers must “gain agency and see themselves as the experts in the room, not whatever AI is telling them.” This administrative concentration mirrors findings from the media sector, where journalists adopted AI for transcription and document search but drew a sharp line at anything that would, in a journalist’s words, “stand in as me.” In education, the line is drawn between acceptable support for preparation and unacceptable substitution for the human relational work of teaching.

Third, in our interviews, the most consequential disruption is occurring not in how teachers deploy AI, but in how students use it, forcing educators into a defensive redefinition of their roles. One founder of an edtech startup, observed that the most heavily funded tools on the market are those that make learning “passive”—automating course video generation, converting textbooks into podcasts, and completing assignments for students. This has shifted the teacher’s function from knowledge facilitator to AI-use monitor. One teacher now requires handwritten, in-class assignments for high school CS students because take-home work is inevitably AI-assisted. As he noted,

*“I’ve observed many students are no longer willing to think independently. When they encounter bugs, their first reaction is to ask AI rather than try a few methods themselves. This dependency worries me.”*

A director of a school research center at Columbia University, described teachers increasingly “AI-proofing” their curriculum—designing tasks that require “productive struggle” and in-person explanation of underlying logic. Rather than integrating AI as a pedagogical asset, several

educators we interviewed are restructuring their instruction to defend against it, a dynamic that speaks to the technology's misalignment with developmental goals rather than its seamless adoption.

## II. Labor Anxieties

As in healthcare and media, workers in education articulate anxieties based on what they can directly observe, not on the full trajectory of AI deployment in their sector. But whereas healthcare workers worry about care quality and journalists about credibility and the training pipeline, teachers' concerns center on the cognitive and relational foundations of education itself.

The most consistent and sector-specific anxiety is the erosion of cognitive autonomy and "productive struggle." Across interviews, educators described a "dependency trap" in which students bypass the difficult intellectual labor that constitutes learning. A Texas third-grade teacher, was explicit:

*"AI is robbing kids from the comprehension ability. We're going to do what makes us comfortable. It's sad kids say they don't want to learn how to write an essay because they have AI."*

An interviewee observed the same pattern in high school computer science: students' first instinct upon encountering a coding error is to consult an AI rather than engage in independent debugging. These teacher observations are reinforced by emerging research on classroom AI use. A recent study from Cambridge University and Microsoft Research found that students who took notes on text passages without AI assistance demonstrated better reading comprehension than those who relied on chatbots for help, suggesting that the cognitive shortcuts AI provides may actively erode learning capacity (NYT, 2026). The Brookings Institution, in a report on school AI use, concluded that for now, "the risks of utilizing AI in education overshadow its benefits". The New York Times documented a Newark high school attempting to address this dependency through a new "AI literacy" course, framed as a "driver's license" for AI, where students are taught to ask: "Are you steering the technology, or is it steering you?" Yet even this well-intentioned response highlights the structural dilemma: schools are now tasked with teaching students how to navigate a technology's cognitive pull, a responsibility that falls unevenly on educators who themselves lack clear guidance.

One senior professor and researcher at Columbia's teacher college emphasized that "productive struggle" is not incidental to learning but constitutive of it—the difficulty itself is where cognitive growth occurs. As one researcher argued, if AI tools are designed to remove that friction, they risk undermining the very human capital formation that education is meant to achieve. This concern is education's parallel to healthcare workers' fear that AI undermines clinical judgment: just as diagnostic algorithms can miss the patient who, in a physician assistant's words, "won't present

with a disease like how it's written in a textbook," educational AI can miss the student whose learning trajectory does not follow the model's predicted path.

A second cluster of anxieties centers on the dehumanization of educational relationships and institutional culture. Teaching is inherently relational work involving emotional regulation, trust-building, and what an expert called the "aha moments"—the point of cognitive breakthrough that only a human observer can recognize and reinforce. One interviewed teacher articulated a fear that AI adoption is eroding these connections at every level of the school ecosystem:

*"Part of the education space is having a connection with one another... It's stripping everybody from their connection."*

Her observation that even internal communications were being drafted by AI signals a cultural shift in which transactional efficiency replaces relational presence. This concern is not about teacher displacement per se—every interviewee agreed that AI cannot replicate the emotional labor of teaching—but about the quiet corrosion of the human environment in which learning occurs. In this line, however, the emotional component of AI also provides reassurance when faced with the question of AI replacement fears. Teachers generally agree that:

*"[...] most of my job is classroom management and interacting with students —AI doesn't really help with that"*

Third, educators raised concerns about exacerbated inequality and, in some contexts, surveillance. An expert in teaching described a stark "two-tiered" classroom experience: students with the latest laptops and GPT-4 subscriptions produce complex projects with AI assistance, while students relying on old school computers or phones can only complete basic functions. As he explained,

*"There's a clear gap between students with the latest laptops who can afford GPT-4 subscriptions, and those using old school computers or even just phones. The former can experiment at home; the latter can only use the computer lab during breaks... Students with resources produce more complex projects because AI helps them solve technical difficulties. Students without resources can only do basic functions. Over time, this gap accumulates."*

This observation is reinforced by quantitative data: the RAND Corporation found that while 67% of low-poverty districts had offered AI training to teachers by Fall 2024, only 39% of high-poverty districts had done so (RAND, 2025). An interview with a teacher from Texas introduced an additional dimension absent from other sectors: in her Texas school district, she feared AI could be weaponized as a surveillance tool in an already politicized educational environment. After a colleague was threatened with firing for social media posts, she worried about who controls AI monitoring systems and whom they serve. This political dimension of AI governance in education—where tools may be deployed to police educators as well as assist them—deserves



attention that the aggregate adoption statistics obscure. This observation is reinforced by quantitative data: the RAND Corporation found that while 67% of low-poverty districts had offered AI training to teachers by Fall 2024, only 39% of high-poverty districts had done so ([RAND](#), 2025).

### **III. Worker Preparedness**

Given the variance in institutional support, educators' preparedness for AI integration ranges from active experimentation to outright denial, with most falling somewhere in between.

First, there is a near-universal absence of meaningful, role-specific training. The training that does exist is often superficial and unequally distributed. Geoffrey Cao described his district's offerings as limited to "how to use ChatGPT," with "little deep training on 'how to redesign curriculum in the AI era.'" And teachers are already busy; adding these trainings causes burnout. What we need is time—time to redesign curriculum, not forced rapid adaptation." Experts noted that prompting and re-prompting AI is a genuine skill, "but not one that is being taught."

Importantly, this training gap tracks a deeper pattern of socioeconomic inequality in school funding. U.S. public schools remain among the most inequitably funded institutions in the developed world, with per-pupil expenditures varying dramatically by district wealth ([Labaree](#), 2008; [Learning Policy Institute](#), 2016). Research consistently shows that these funding disparities translate into gaps in instructional materials, teacher quality, and now—critically—access to technological capacity-building. The RAND Corporation's American School District Panel confirms this pattern specifically for AI training: while 67% of low-poverty districts had offered AI training to their teachers by Fall 2024, only 39% of high-poverty districts had done so ([RAND](#), 2025). A teacher interviewee from Texas, Anna Clarke, shared the severe economic constraints her school was facing—"the other day we had a pizza party fundraiser to afford buying paper in the school"—yet when asked about AI training, she denied having had any such opportunity. The result is a workforce expected to adapt to a transformative technology without the time, resources, or pedagogical guidance to do so.

Second, preparedness is constrained by ecological fragmentation. Effective AI integration requires alignment across district leaders, school administrators, and teachers—a systems-level coherence that decentralized governance undermines. Several experts interviewed emphasized an "ecological approach" in which all three stakeholder groups share a vision for AI use. Yet in practice, district-level decentralization, particularly in large urban systems like New York City, means that AI policies and resources vary dramatically from school to school. A vocational training instructor, encountered a parallel problem: corporate software systems update quarterly, but curriculum approval cycles lag far behind, leaving students to learn specific interfaces on the job rather than in the classroom. As he noted,

*“We cooperate with several companies; their software systems update every quarter. Our curriculum syllabus was approved last year; by the time it’s approved, the interface has changed. We can only teach students basic principles; specific interface operations they learn on the job.”*

This temporal misalignment between institutional processes and technological change means that educators are often working with outdated guidance relative to the tools their students will encounter in the workplace, a gap that is particularly consequential in vocational and career-training contexts.

Third, teachers are structurally isolated from collective knowledge and voice infrastructure. A teacher’s statement that “the only place I have power is inside my classroom” encapsulates the atomization of educator labor in the face of AI. Unlike media and entertainment, where unions such as the NewsGuild-CWA and SAG-AFTRA have actively bargained AI protections, teacher unions’ engagement on AI has been comparatively less visible in our research.

Similarly, the American Federation of Teachers, representing 1.8 million members, has taken a multi-pronged approach. In July 2025, the AFT—in partnership with its New York City affiliate, the United Federation of Teachers—launched the National Academy for AI Instruction, a \$23 million initiative backed by Microsoft, OpenAI, and Anthropic, with a five-year goal of training 400,000 educators (AFT, 2025). The Academy offers workshops, credential pathways, and a mechanism for teachers to provide feedback to AI developers. AFT President has emphasized that the union is “concerned about safety, and we’re really concerned about people and kids’ thinking,” noting explicitly, “We’re not worried about being replaced” (Nextgov, 2025). The AFT has also passed resolutions declaring AI impact a mandatory subject of collective bargaining and committing to develop contract language protecting workers from displacement (AFT, 2024).

Yet these union efforts have barely penetrated the day-to-day awareness of the frontline educators we interviewed. By comparison, education, training, and library occupations have among the highest union membership rates of any occupational group—32.3 percent nationally in 2024, second only to protective service occupations (BLS, 2025). The disconnect between this relatively high organizational density and the limited awareness of AI-specific union initiatives suggests that even where unions are active, their AI work operates as an exceptional site of policy-making that has not yet translated into widespread worker-level engagement. Aphilanthropic program officer captured this gap when she noted that “a core principle being upheld is that the teacher is not going to be replaced”—a defensive posture that may leave more complex questions about working conditions and pedagogical autonomy unaddressed. Without cross-district networks to share “micro-innovations” or to engage AI vendors from an informed, collective position, teachers are left to navigate deployment decisions alone. This isolation is itself a form of unpreparedness: workers cannot advocate for or against changes they have no forum to discuss.

## IV. Sector Interventions

### 1. Rebalance AI-earmarked school resources to mitigate inequalities in its adoption

Following a Google.org’s AI and education research grantee’s framework, an ecological approach is best suited to meet the challenge. Poorer school districts should be prioritized in funding for AI capacity-building programs and expert appointments to ensure teachers who are already under severe pressure –and ultimately their students– are not left behind. This includes multi-stakeholder alignment initiatives that bring together district leaders and officials with principals and teachers to ensure programming is co-designed rather than imposed top-down.

In addition, expert staffing with AI coordinators should provide guidance to teachers by offering practical examples and opportunities to reflect beyond ChatGPT usage workshops to maximize the awareness of the need to preserve “productive struggle”. Philanthropic capital can help bridge the current AI training capacity gap between richer and poorer districts by both supplementing these efforts and convening the education community advocating for long-term public investment that tackles the systemic issues with school district funding. In this line, the Humanity AI coalition’s education line of work, where the Omidyar Network is already involved, could play a convening role to push more equitable funding for AI adaptation across school districts ([Humanity AI](#), 2026).

### 2. Fund net-positive AI initiatives in education showcasing worker-centered design

While much of the media focus has been centred around edtech startups providing self-contained products to gamify and digitize the teaching of some subjects, these can also generate backlash and reluctancies due to the exaggerated techno-optimistic narratives they push. However, finding and supporting meaningful implementations for AI can contribute to change the narrative through examples such as the following:

- Playlab: Enables teachers to build their own bots and apps with embedded pedagogical rigor, positioning educators as designers rather than consumers.
- TeachingLab Studio: Focuses on high-quality professional learning experiences that integrate AI literacy with content-specific pedagogy.
- Quill: Develops open-source writing feedback tools designed with educator input to augment rather than automate formative assessment.
- Decoded Futures: Takes a systems approach to supporting districts in developing comprehensive AI strategies aligned with equity goals, such as their case study on Teach for America.

Interviewed philanthropic leaders in the education space emphasize that "centring instructional rigor is sometimes neglected versus the shiny new tool." Philanthropic investment should reward initiatives that demonstrate measurable improvements in teacher autonomy, student engagement, and real learning outcomes.

### **3. Support worker voice and collective action infrastructure**

The interviews reveal that teachers feel isolated in their responses to AI, with individual educators retreating to "the only place I have power, inside my classroom." In fact, when asked about potential spaces to seek collective solutions some mentioned not knowing where to go. This atomization prevents knowledge exchange and leaves workers vulnerable to top-down mandates and uncertainty. In this line, solutions might include supporting labor union capacity to develop AI policies and contract language that protect workers while enabling beneficial uses. This includes supporting technical expertise within unions so they can negotiate from informed positions.

Similarly, organizing cross-district teacher networks that enable educators to share "micro-innovations" about what works and what doesn't, creating peer-driven knowledge rather than relying solely on vendor claims or administrative directives. As an expert from Teacher's College at Columbia notes, "we need to involve AI companies who are ready to restrict their own development in conversation with the education sector." Philanthropic funding can create neutral spaces for these conversations where worker perspectives carry equal weight. However, learnings should be found in past experiences of philanthropic-driven advocacy efforts seeing mixed results. This is the case of the Intensive Partnerships for Effective Teaching program by the Gates Foundation ([RAND](#), 2018), in an attempt to link teacher pay to student test results.



## Recommendations

The Omidyar Network is well-positioned to address the systemic challenges of AI through its integrated focus on policy change and worker voice. By leveraging the Workers/Jobs, Media, and Coalitions programs grantmaking resources, the organization can move toward a proactive model of technological governance. These interventions should be concentrated into 3 strategic pillars: (1) knowledge sharing, (2) regulatory advocacy and (3) worker co-design.

Interventions are organized by feasibility, with **green** denoting readily implementable actions, **orange** indicating moderate complexity, and **red** representing longer-term efforts.

## Intervention #1: Knowledge Sharing & Transparency

The gap between institutional deployment of AI tools and worker knowledge is one of the most consequential barriers to worker empowerment. Thus, we believe knowledge sharing is a high-leverage, low-cost intervention point.

- **Sector-Specific Research Centers and Fellowships:** Despite growing academic and market research, much of it analyzes macro-level impact without centering workers themselves. Sector-specific research centers that embed researchers within the industries they study produce qualitatively different insights. For instance, University of Toronto's [Creative Labour Critical Futures](#) (CLCF) project and the [DigiLabour](#) research hub are strong examples. Both centers bring together faculty across arts, culture, and media fields to produce original research on worker-led governance. Omidyar can fund similar centers for other sectors or support fellowships to ensure worker-centered research.
- **Shared Library of Bargaining Contracts:** Unions that have negotiated AI protections have produced some of the most significant labor language, but they are scattered across websites, difficult to compare, and invisible to the majority of non-union workers. A searchable database of AI contract clauses, organized by sector and by issue, would help unions benchmark new negotiations, give non-union workers visibility into existing labor protections, and accelerate cross-sector learning. While the University of Toronto's CLCF maintains a global [tracker](#) of worker mobilizations around AI in media and entertainment, no database exists for negotiated contracts.
- **Narratives Centering Workers Voices:** Funding journalism that center workers' direct experiences—including their productivity gains, frustrations, and creative workarounds—gives visibility to perspectives the mainstream narrative excludes. Brookings Metro's [Workers and AI project](#), launched in 2024, pursues exactly this: a storytelling series centering worker experiences alongside quantitative analysis and cross-sector convenings. Omidyar can fund similar initiatives through existing journalism nonprofits or its Reporters in Residence program.
- **Employer Training:** Employer-led training programs or fellowships can equip workers to use AI responsibly and effectively by teaching practical skills, like prompting and evaluating outputs, while establishing clear boundaries on appropriate use. One of our interviewees described a company-wide training at an American film studio that emphasized human oversight in sensitive domains (like hiring decisions and performance evaluations). The training also covered how to use AI responsibly and improve prompting tailored to employees' workflows. Omidyar can help fund, standardize, and scale employer-led training that might otherwise be under-resourced.



## Intervention #2: Regulatory Advocacy

There is a growing demand for policy approaches that incentivize human capital and support workers through periods of displacement. For the Omidyar Network's lobbying efforts, some promising avenues include incentivizing employers through tax credit schemes and supporting reemployment for workers affected by AI-driven change.

- **Tax-Based Incentives**

- **Workforce Training Tax Credits:** One policy avenue that warrants additional exploration is a federal [AI Workforce Training Tax Credit](#), which would allow employers to claim a 30% tax credit on qualified AI training expenditures, capped at \$2,500 per employee per year, covering accredited courses, internal instruction, certificate programs, and workshops in areas like prompt engineering, AI ethics, and data literacy. The underlying policy logic is to shift employer incentives toward sustained investment in human capital as AI adoption accelerates across industries.
- **Wage & Hiring Tax Credits:** Similar incentives can be applied for human capital. New York's wage and hiring tax credit [program](#) for media industries links public subsidies to employee retention and the creation of new full-time jobs through structured tax incentives. The Empire State Newspaper and Broadcast Media Jobs Program, provides a credit equal to 50% of eligible employee wages for retention, alongside a \$5,000 credit for each full-time job created year-over-year. The program is intended to encourage workforce stability and job creation, particularly for smaller employers, against AI-driven changes.
- **Safety Net Insurance:** Efforts to adapt safety net insurance to an AI-driven labor market are emerging at the state level through reforms that strengthen advance notice requirements and embed transition support prior to layoffs. Bills in [New York](#), [Minnesota](#), and [California](#) collectively strengthen advance notice by requiring a 90 day notice for layoffs affecting 25+ employees and embedding transition support prior to layoffs. These changes strengthen early warning systems that give workers more time to access unemployment benefits or seek new employment before displacement.

### Intervention #3: Worker Co-Design

- **Sector-Specific Worker Boards for AI Oversight:** Sector-specific worker boards provide a formal mechanism for embedding worker voice across the full lifecycle of AI deployment, including governance, use case selection, and implementation. In Pennsylvania, state government employees represented by SEIU Local 668 secured a contract establishing a [worker board](#) to oversee generative AI tools in public sector workflow, reflecting ongoing worker inputs at the workplace level.
- **Worker Experimentation Incentives:** Small grants for workers to test, evaluate, and develop AI tools suited to their workflows would generate ground-up intelligence on what works for workers. The [IPI Global AI Accelerator](#), funded by the Patrick J. McGovern Foundation, provides news outlets across the Global South with funding, training, and advisory support to develop AI tools tailored to their contexts. While oriented toward newsrooms, its bottom-up structure shows that investing in context-specific experimentation produces more useful outcomes. Adapting this model to other sectors by putting experimentation in workers' hands, could help address the gap our research identifies.
- **Labor-Technology Partnerships for Co-Design:** Labor-technology partnerships, like the [partnership](#) between The American Federation of Labor and Congress of Industrial Organizations and Microsoft, demonstrate how worker input can be embedded upstream in AI development. The collaboration focuses on sharing technical knowledge with workers, integrating labor expertise into system design, and shaping public policy around workforce impacts – a structured channel for aligning innovation with labor needs.
- **Worker-Centered AI Tools:** The National Domestic Workers Alliance's ["Ask Aya"](#) demonstrates how AI can be intentionally designed to strengthen worker outcomes. Designed to support domestic workers in navigating workplace negotiations, the tool provides real-time, context-specific guidance that translates collective expertise into accessible formats. Its strong engagement reflects the value of AI systems that are practical, responsive, and directly aligned with worker needs, particularly in sectors with limited formal protections.

# Conclusion

This report set out to examine three core questions: to what extent media narratives about AI and work reflect what is happening in workplaces; how AI is changing day-to-day workplace experiences, and how workers are responding; and where there are opportunities for philanthropic and policy interventions that support worker agency. Across four sectors and three methodological lenses, including media analysis, worker surveys, and qualitative interviews, a consistent gap emerged within the scope of this research. Public narratives often center on rapid, widespread job displacement. In contrast, the patterns we observed point to a more fragmented, uneven, and often less visible process of AI integration into everyday work.

Workers in our sample were not passive. Many described adapting to new tools, negotiating how those tools were used, and, in some cases, pushing back, largely through informal channels and often without adequate institutional support. The impacts they identified were not primarily mass layoffs, but more diffuse shifts like erosion of entry-level pipelines, changes in professional autonomy, shifts in how learning and judgment are exercised, and unresolved questions around consent and compensation when worker-generated data is used to train AI systems. Across sectors, our research surfaced a recurring set of conditions: AI tools introduced without advance notice or training; formal policies that existed on paper but were not widely known or consistently enforced; labor protections that applied unevenly, particularly between unionized and non-union workers; and public narratives that amplified executive and expert perspectives more than worker experience. These patterns were consistent within our dataset, though the limitations below caution against treating them as universal.

The philanthropic interventions proposed in this report, including sector-specific research centers, shared bargaining databases, worker co-design boards, and AI transparency standards, are grounded in these findings. They are designed to respond to the structural gaps our research surfaces, including limited worker voice in deployment decisions, information asymmetries between management and labor, and the constraints of existing labor protections for workers most exposed to AI-driven change.

## Limitations

This research has a few limitations that directly shape how its findings should be interpreted.

Our survey of over 80 workers surfaces consistent patterns and directional insights, but it is not statistically representative of the broader U.S. workforce and should not be read as such. Respondents were drawn primarily from the New York and New Jersey metropolitan area, which shapes the sample in important ways. Workers in this region are more likely to operate in unionized or professionally networked environments, to have encountered public discourse about AI and labor, and to have institutional or social venues for discussing those concerns. Workers in other geographic contexts including regions with lower union density, different economic structures, or less exposure to technology sector narratives may hold different perceptions of risk, agency, and institutional trust.

The survey also skewed toward workers with some connection to professional networks and, in some cases, graduate or professional school settings, environments where critical reflection on workplace technology may be more normalized. This likely influenced both who responded and how concerns were articulated. Workers who are more isolated from these networks, more exposed to retaliation risk, or less engaged with public AI discourse are less represented in our findings, despite potentially facing more acute or less visible forms of disruption.

A related limitation applies to our interview sample. While the project was designed to include administrative, entry-level, and lower-level workers, these groups were more difficult to reach through our recruitment channels. The sample also includes some mid-level workers and organizers engaged on these issues. As a result, the interview pool skews toward individuals with greater institutional visibility, including those connected to unions, professional associations, or academic-adjacent environments. This means administrative and entry-level workers are underrepresented in our qualitative findings relative to their presence in the broader workforce.

Our findings are most appropriately interpreted within the scope of these research questions. They are not designed to support broader claims about AI's overall economic impact or to predict displacement at scale.

## **Areas for Further Research**

Several questions raised by this research remain unresolved and warrant further investigation. First, the data infrastructure for understanding AI's labor market impact remains underdeveloped, a constraint this project encountered directly. Federal data collection does not currently track how AI is reshaping employment at the task, occupation, or firm level, making it difficult to distinguish structural transformation from cyclical disruption or narrative-driven "AI-washing." The proposed [Workforce Transparency Act](#), introduced by Senators Mark Warner and Ted Budd, would establish a federal framework for collecting aggregated, anonymized data on AI usage, including task-level applications, geographic distribution, and changes over time. This type of infrastructure is a prerequisite for evidence-based policy. Philanthropy may play a role both in supporting such frameworks and in complementing them with worker-centered qualitative research that captures what aggregate data cannot.

Second, how productivity gains from AI are distributed remains largely absent from current policy discourse. If AI generates economic value, the mechanisms through which that value reaches workers, rather than accruing primarily to capital, require further analysis. Universal Basic Income has re-entered public debate as one possible mechanism, but existing research on its labor market effects is mixed, and questions of fiscal feasibility and political viability remain unresolved. Alternative approaches, including sectoral profit-sharing arrangements or expanded investment in retraining and transition support, also require evaluation. This report is not positioned to determine which approach is most viable, but it highlights the urgency of the question as entry-level pathways show signs of narrowing in our findings.

Third, both this project and the broader research landscape are geographically and demographically concentrated. Future research should deliberately recruit workers outside major metropolitan areas, in non-union employment settings, and in lower-wage service and care roles, where AI-driven change may be unfolding with less visibility and fewer protective structures. Our own difficulty reaching entry-level and administrative workers points to a methodological challenge that future research should address directly.

Finally, the pace of change presents an additional research challenge. Several dynamics identified in this report, including AI integration into content production pipelines, forms of algorithmic management, and changes to training pathways, were still emerging or accelerating at the time of writing. Longitudinal research that tracks worker experience over time will be necessary to distinguish short-term disruption from longer-term structural change and to evaluate whether the interventions proposed here are having their intended effects.

Within our sample, workers are already making decisions about how to engage with these systems, including which tools to adopt, which concerns to raise, and which risks to absorb, often without consistent access to information or institutional support. The gap between what workers described needing and what was available to them is the most consistent finding across this research and the clearest rationale for the interventions proposed in this report.

# References

- Acemoglu, D., & Restrepo, P. (2018). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6), 2188–2244.
- AFT (2024). Resolution on Artificial Intelligence. <https://www.aft.org/resolution/artificial-intelligence>
- AFT (2025). AFT to Launch National Academy for AI Instruction with Microsoft, OpenAI, Anthropic, and United Federation of Teachers. <https://www.aft.org/press-release/aft-launch-national-academy-ai-instruction-microsoft-openai-anthropic-and-united>
- Apell, P., Locher, S., Milde, A., Eriksson, H., & Eriksson, H. (2026). Explaining the slow adoption of AI innovations in health care: Network analysis approach. *JMIR AI*, 5, e60458. <https://doi.org/10.2196/60458>
- Axios. (2024). Axios and OpenAI expand partnership to scale local news [Press release]. <https://www.axios.com/press-past-releases/axios-openai-expand-partnership-to-scale-local-news>
- Baker, B. D., Farrie, D., & Sciarra, D. G. (2020). Closing America’s education funding gaps. The Century Foundation. <https://tcf.org/content/report/closing-americas-education-funding-gaps/>
- Bracken, A., Reilly, C., Feeley, A., Sheehan, E., & Merghani, K. (2025). Artificial intelligence (AI)-powered documentation systems in healthcare: A systematic review. *Journal of Medical Systems*, 49(28). <https://doi.org/10.1007/s10916-025-02157-4>
- Brynjolfsson, E., & McAfee, A. (2014). The second machine age: Work, progress, and prosperity in a time of brilliant technologies. W. W. Norton & Company.
- Chen, W. X., Srinivasan, S., & Zakerinia, S. (2024). Displacement or complementarity? The labor market impact of generative AI. Harvard Business School Working Paper No. 25-039. <https://www.hbs.edu/faculty/Pages/item.aspx?num=67045>
- Dominski, J., & Lee, Y. S. (2025). Advancing AI capabilities and evolving labor outcomes (arXiv:2507.08244). arXiv. <https://arxiv.org/abs/2507.08244>
- Duke Health AI Evaluation & Governance Program. (2021). Algorithm-Based Clinical Decision Support (ABCDs) Oversight. <https://healthaigovernance.duke.edu/abcs-oversight>
- EdWeek (2024). NEA Approves AI Guidance, But It's Vital for Educators to Tread Carefully (citing Center for Democracy & Technology survey data: 71% of educators received no AI professional development). <https://www.edweek.org/teaching-learning/nea-approves-ai-guidance-but-its-vital-for-educators-to-tread-carefully/2024/07>
- Forrester. (2026, January 13). The Forrester AI job impact forecast, U.S., 2025–2030 [Press release]. <https://www.forrester.com/press-newsroom/forrester-impact-ai-jobs-forecast/>
- Fortune. (2026, February 20). New York nurses union raises AI safeguards in deal after longest strike. <https://fortune.com/2026/02/20/new-york-nurses-union-raise-ai-safeguards-deal-longest-strike/>
- Fox Business. (2026). Nvidia announces AI-powered healthcare agents outperform nurses, cost \$9/hour. <https://www.foxbusiness.com/technology/nvidia-a-announces-ai-powered-health-care-agents-outperform-nurses-cost-9-hour>
- Grieco, E. (2022, August 4). About one-in-six U.S. journalists at news outlets are part of a union; many more would join one if they could. Pew Research Center. <https://www.pewresearch.org/short-reads/2022/08/04/about-one-in-six-u-s-journalists-at-news-outlets-are-part-of-a-union-many-more-would-join-one-if-they-could/>



- Grieco, E. (2023, April 4). U.S. journalists' beats vary widely by gender and other factors. Pew Research Center. <https://www.pewresearch.org/short-reads/2023/04/04/us-journalists-beats-vary-widely-by-gender-and-other-factors/>
- Griner, D. (2023, December 27). The Times sues OpenAI and Microsoft over A.I. use of copyrighted work. The New York Times. <https://www.nytimes.com/2023/12/27/business/media/new-york-times-open-ai-microsoft-lawsuit.html>
- Gupta, A. (2026, April 6). Associated Press starts offering buyouts to newspaper journalists amid wider AI transformation. Fortune. <https://fortune.com/2026/04/06/associated-press-starts-offering-buyouts-to-newspaper-journalists-amid-wider-ai-transformation/>
- Hackensack Meridian Health. (2022). Artificial intelligence (AI) governance. <https://www.hackensackmeridianhealth.org/en/about-us/ai/ai-governance>
- Kelley, L. (2026, February 1). Layoffs and AI-washing. The New York Times. <https://www.nytimes.com/2026/02/01/business/layoffs-ai-washing.html>
- Labaree, D. F. (2008). The winning ways of a losing strategy: Educationalizing social problems in the United States. *Educational Theory*, 58(4), 447–460. <https://www.jstor.org/stable/41238070>
- Learning Policy Institute. (2016). How money matters: Education funding and student outcomes. <https://learningpolicyinstitute.org/product/how-money-matters-factsheet>
- Mahar, S., Meyer, B., & Parker, N. (2026). Artificial intelligence, productivity, and the workforce: Evidence from corporate executives (Working Paper No. 2026-4). Federal Reserve Bank of Atlanta. <https://www.atlantafed.org/research-and-data/publications/working-papers/2026/03/25/04-artificial-intelligence-productivity-and-the-workforce-evidence-from-corporate-executives>
- Maslej, N., et al. (2025). Artificial intelligence index report 2025. Stanford Institute for Human-Centered AI. <https://hai.stanford.edu/ai-index/2025-ai-index-report>
- McCarty Carino, M. (2025, September 9). AI is costing jobs, but not always the way you think. Marketplace. <https://www.marketplace.org/story/2025/09/09/ai-is-costing-jobs-but-not-always-the-way-you-think>
- McCarty Carino, M. (2026, February 3). Recent layoffs have been attributed to AI. Is that reality? Marketplace. <https://www.marketplace.org/story/2026/02/02/recent-layoffs-have-been-attributed-to-ai-is-that-reality>
- McKinsey Global Institute. (2023). The economic potential of generative AI: The next productivity frontier. McKinsey & Company.
- Menlo Ventures. (2025, October 21). The state of AI in healthcare. <https://menlovc.com/perspective/2025-the-state-of-ai-in-healthcare/>
- Menlo Ventures. (2025, December 9). The state of generative AI in the enterprise. <https://menlovc.com/perspective/2025-the-state-of-generative-ai-in-the-enterprise/>
- Microsoft. (n.d.). Nuance partner professional services: Azure Machine Learning customer story. <https://www.microsoft.com/en/customers/story/1415193159282858383-nuance-partner-professional-services-azure-machine-learning>
- Microsoft. (2026, January). AI in education report: Insights to support teaching and learning. <https://www.microsoft.com/en-us/education/blog/2026/01/ai-in-education-report-insights-to-support-teaching-and-learning>
- National Nurses United. (2024). AI's impact on nursing and health care. <https://www.nationalnursesunited.org/artificial-intelligence>

- National Nurses United. (2024, May 15). National nurses United survey finds AI technology degrades and undermines patient safety [Press release]. <https://www.nationalnursesunited.org/press/national-nurses-united-survey-finds-ai-technology-undermines-patient-safety>
- NBC News. (2025). ChatGPT health: Under-triaged half medical emergencies. <https://www.nbcnews.com/health/health-news/chatgpt-health-under-triaged-half-medical-emergencies-rcna261409>
- Nanevski, I., Jäger, S., Schulte-Althoff, M., Behnke, E.-M., Fürstenau, D., & Biessmann, F. (2025). The potential of AI in nursing care: Multicenter evaluation in fall risk assessment. *Journal of Medical Internet Research*, 27, e71034. <https://doi.org/10.2196/71034>
- Newman, N. (2026). Journalism, media and technology trends and predictions 2026. Reuters Institute for the Study of Journalism. <https://reutersinstitute.politics.ox.ac.uk/journalism-media-and-technology-trends-and-predictions-2026>
- Nextgov (2025). Teachers Union Looks to States and Industry for AI Regulation Absent Federal Action. <https://www.nextgov.com/artificial-intelligence/2025/08/teachers-union-looks-states-and-industry-ai-regulation-absent-federal-action/407375/>
- Northwestern University Local News Initiative. (2025). The state of local news 2025. <https://localnewsinitiative.northwestern.edu/projects/state-of-local-news/2025/report/>
- Pailaha, A. D. (2023). The impact and issues of artificial intelligence in nursing science and healthcare settings. *SAGE Open Nursing*, 9. <https://doi.org/10.1177/23779608231196847>
- Palantir Technologies. (n.d.). Palantir for hospitals. <https://www.palantir.com/offerings/palantir-for-hospitals/>
- Pew Research Center. (2022). About one-in-six U.S. journalists at news outlets are part of a union; many more would join one if they could. <https://www.pewresearch.org/short-reads/2022/08/04/about-one-in-six-u-s-journalists-at-news-outlets-are-part-of-a-union-many-more-would-join-one-if-they-could/>
- Pew Research Center. (2023). U.S. journalists' beats vary widely by gender and other factors. <https://www.pewresearch.org/short-reads/2023/04/04/us-journalists-beats-vary-widely-by-gender-and-other-factors/>
- Platform Cooperativism Consortium. (2025). AI Labor Council for Healthcare: Exploring Models for Worker-led AI Governance. <https://platform.coop/blog/16684/>
- Reuters Institute for the Study of Journalism. (n.d.). AI and journalism: What's next. University of Oxford. <https://reutersinstitute.politics.ox.ac.uk/news/ai-and-journalism-whats-next>
- Robertson, K. (2025, November 7). A.I. is sweeping through newsrooms. *The New York Times*. <https://www.nytimes.com/2025/11/07/business/media/ai-news-media.html>
- Scire, S. (2025, September). In France, AI revenue is going directly to journalists. Could that happen in the U.S.? Nieman Lab. <https://www.niemanlab.org/2025/09/in-france-ai-revenue-is-going-directly-to-journalists-could-that-happen-in-the-u-s/>
- Shumer, M. [@mattshumer\_]. (2026, February 10). Something big is happening. [Tweet]. X. [https://x.com/mattshumer\\_/status/2021256989876109403](https://x.com/mattshumer_/status/2021256989876109403)
- Singla, A., Sukharevsky, A., Yee, L., Chui, M., & Hall, B. (2025). The state of AI: How organizations are rewiring to capture value. McKinsey & Company.
- Stecher, B. M., Holtzman, D. J., Garet, M. S., Hamilton, L. S., Engberg, J., Steiner, E. D., ... & Chambers, J. (2018). *Improving teaching effectiveness: Final report: The Intensive Partnerships for Effective Teaching through 2015–2016*. RAND Corporation.

[https://www.rand.org/pubs/research\\_reports/RR2242.html](https://www.rand.org/pubs/research_reports/RR2242.html)

Takita, H., Kabata, D., Walston, S. L., Tatekawa, H., Saito, K., Tsujimoto, Y., Miki, Y., & Ueda, D. (2025). A systematic review and meta-analysis of diagnostic performance comparison between generative AI and physicians. *npj Digital Medicine*, 8(1), 175.  
<https://doi.org/10.1038/s41746-025-01543-z>

UCLA Health. (2022). UCLA Health AI Council (HAIC).  
<https://www.uclahealth.org/health-ai/ucla-health-ai-council-haic>

UCLA Health. (2026, February 27). AI tools are reducing burnout and transforming patient-centered care.  
<https://www.uclahealth.org/news/publication/ai-tools-are-reducing-burnout-and-transforming-patient>

U.S. Congress. (2023). NO FAKES Act of 2023: Nurture originals, foster art, and keep entertainment safe act (118th Congress).

UNESCO. (2023). Guidance for generative AI in education and research.  
<https://unesdoc.unesco.org/ark:/48223/pf0000386693>

Warner, M. R., & Budd, T. (2026, April 30). Warner, Budd introduce legislation to collect data on AI's impact, guide lawmakers. U.S. Senate.  
<https://www.warner.senate.gov/newsroom/pre>

[ss-releases/warner-budd-introduce-legislation-to-collect-data-on-ai-impact-guide-lawmakers/](https://www.warner.senate.gov/newsroom/pre/ss-releases/warner-budd-introduce-legislation-to-collect-data-on-ai-impact-guide-lawmakers/)

Wells, K. J., & Ustek Spilda, F. (2024). Uber for nursing: How an AI-powered gig model is threatening health care. Roosevelt Institute.  
<https://rooseveltinstitute.org/publications/uber-for-nursing/>

Wells, K. J., & Ustek Spilda, F. (2025). Uber for nursing part II. AI Now Institute.  
<https://ainowinstitute.org/publications/uber-for-nursing>

World Economic Forum. (2025). The future of jobs report 2025.  
<https://www.weforum.org/reports/>

Ye, X. M., & Ranganathan, A. (2026). AI promised to free up workers' time. UC Berkeley Haas researchers found the opposite. UC Berkeley Haas Newsroom.  
<https://newsroom.haas.berkeley.edu/ai-promised-to-free-up-workers-time-uc-berkeley-haas-researchers-found-the-opposite/>

Zhao, L., Liu, S., Xin, T., Tan, J., Wang, X., Li, Y., Bian, Z., Chen, Y., Kong, F., Bian, J., Qian, C., & Zhang, Z. (2026). AI agent in healthcare: Applications, evaluations, and future directions. *npj Artificial Intelligence*, 2(31).  
<https://doi.org/10.1038/s44387-026-00076-4>

# Appendix

## A. Media Narrative Analysis Methodology

Our media analysis maps the ecosystem of narratives shaping public and policy debates about AI's impact on work. The goal was to understand how different types of publications, from academic research centers to independent tech newsletters, frame the relationship between artificial intelligence and worker displacement. Below is a brief overview of the coding categories we created:

- **Industry:** We coded each article by industry sector. Our primary sectors of interest are Health, Entertainment, Media, and Education, though additional sectors were included when articles were substantively relevant to worker displacement themes. We then identified which worker group, if any, was the actual subject of the article, distinguishing between entry-level & clerical workers (clerks, secretaries, admin assistants, receptionists, call center agents), junior/entry-level workers (recent graduates, early-career hires), white collar broadly (used when coverage addressed office workers without specifying role level), professional workers (lawyers, doctors, credentialed roles), and blue collar/manufacturing workers.
- **Worker Voice:** We coded each article Yes/No: does the journalist actually interview an entry-level worker, or is the article based solely on CEO statements, expert analysis, or institutional data?
- **AI Narrative:** We employed a three-category typology to describe AI's relationship to work, consisting of replacement (framed as a direct substitute for a human worker), augmentation (framed as a tool that makes workers more productive while keeping them in charge), and surveillance (used primarily to monitor worker productivity or compliance rather than to replace or assist)
- **Narrative Tone:** We used a five-point scale to assess the overall editorial tone of each article (see below).
- **Voices & Sources Cited:** We tracked whose voice appeared in each article by coding the primary sources quoted or cited. Categories included: Tech CEO or executive; Economist or academic researcher; Affected worker/employee; Labor union representative; Government official or policymaker; HR professional or recruiter; Journalist (analysis piece); and "no human sources quoted."
- **Proposes a Solution:** A simple Yes/No field noting whether the article proposes concrete policy, institutional, or individual responses.

When designing the coding categories for narrative tone, we used a five-point scale to assess the overall editorial tone of each article:

| Score | Label                  | Description                                                                                        |
|-------|------------------------|----------------------------------------------------------------------------------------------------|
| 1     | Alarmist               | AI portrayed as an imminent, large-scale threat using sweeping language with few caveats.          |
| 2     | Concerned but Measured | Acknowledges meaningful displacement risk; balances with data, policy proposals, or expert debate. |

| Score | Label                | Description                                                                                      |
|-------|----------------------|--------------------------------------------------------------------------------------------------|
| 3     | Neutral / Factual    | Reports events or statistics without editorial framing; lets readers draw their own conclusions. |
| 4     | Optimistic           | Frames AI as a net positive: creating new jobs, augmenting productivity, or empowering workers.  |
| 5     | Dismissive / Denying | Actively downplays or rejects the premise that AI poses significant displacement risk.           |
| Mixed | Ambiguous            | Genuinely contradictory frames present; no dominant tone after careful reading.                  |

We chose to look at articles between August 2025-March 2026, capturing the period when agentic AI tools became widely discussed and when Brynjolfsson et al.'s (2025) initial empirical evidence of entry-level employment impacts began circulating. We selected 14 outlets across five categories, deliberately seeking sources that operate at different points in the agenda-setting chain:

| Category                    | Outlets                                                                          |
|-----------------------------|----------------------------------------------------------------------------------|
| Mass Media & National Press | The New York Times, Washington Post, The Guardian (US), Bloomberg, The Economist |
| Digital News & Commentary   | Vox, Slate, The Atlantic                                                         |
| Insider Tech Platforms      | Stratechery, Platformer                                                          |
| Think Tank & Institutional  | McKinsey, Brookings, The World Economic Forum                                    |
| Academic Research Centers   | Stanford Human-Centered AI, Stanford Digital Economy Lab                         |

Team members were each assigned a subset of the 14 outlets. For each outlet, researchers searched for headlines using the outlet's native search bar and applied systematic keyword pairs combining an AI term (e.g., 'artificial intelligence,' 'AI agents,' 'automation') with a labor term (e.g., 'job loss,' 'workforce reduction,' 'future of work'). We then collected all articles within the relevant time period, coding each article independently according to the coding categories. Articles were then peer-reviewed by a second team member to ensure inter-coder reliability on contested or ambiguous cases.

## B. Survey Template

We deployed a survey via SurveyMonkey to assess respondents' backgrounds, workplace experience with AI, and suggested interventions.

## Your Background

1. **What industry do you currently work in?**
  - Media / Journalism
  - Entertainment (film, TV, music)
  - Healthcare
  - Education
  - Other: \_\_\_\_
2. **What is your level of work experience?**
  - None
  - Less than 1 year
  - 1-5 years
  - 6-10 years
  - 10+ years
3. **Are you a member of a union or professional association?**
  - Yes
  - No
  - Unsure
4. **Do you work**
  - In person
  - Remote/virtual
  - Mixed
5. **Is your role primarily**
  - External (e.g. customer facing)
  - Internal
  - Mixed
6. **Which best describes your role type?**
  - Admin/Support role (e.g., clerical, assistant roles)
  - Entry-level (Junior/Intern)
  - Mid-level (Associate/Manager/Resident)
  - Senior-level (Director/VP/Attending)
  - Executive (C-suite)
  - Not applicable

## Your AI Experience

- **How has AI or automation changed your workload?**
  - Significantly decreased
  - Somewhat decreased
  - No change
  - Somewhat increased
  - Significantly increased

- **How do you feel about AI in your workplace?**
  - Very optimistic
  - Somewhat optimistic
  - Neutral
  - Somewhat worried
  - Very worried
- **If you do feel concerned about AI impacting your job or career, what has most shaped that concern? (Select all that apply)**
  - Social Media
  - News Articles
  - AI tools in my workplace
  - Co-workers
  - Friends/Family
  - I'm not concerned about AI impacting my work
- **To what extent is AI coverage in the news and social media causing you anxiety?**
  - 1 - Not at all
  - 2
  - 3
  - 4
  - 5 - Very anxious
- **How accurate do you find news coverage of AI's impact on your field?**
  - Very inaccurate
  - Somewhat inaccurate
  - Neutral
  - Somewhat accurate
  - Very accurate

## Your Workplace

- **Has your employer introduced AI tools or automation in the past two years?**
  - Yes, many
  - Yes, few
  - No
  - Not sure
- **Which of the following has happened in your workplace as a result of AI? (select all that apply)**
  - Reduction in team size
  - Changes in job descriptions
  - Required AI integration into workflows

- Tasks reassigned to AI tools
- None of the above
- **How much of your current daily work do YOU think could be done by AI?**
  - 0-10%
  - 10-40%
  - 40-60%
  - 60-80%
  - 80-100%
  - Unsure

### AI Adoption and Inclusion

- **Have YOU ever been asked for input on how AI or automation should be used in your role?**
  - Yes
  - No
- **How satisfied are you with the amount of input you've had?**
  - I wish I were consulted more (about how AI/automation should be used)
  - I've been consulted the right amount
  - I have been consulted more than needed
- **Have you challenged how AI/automation is or will be used in your workplace?**
  - Yes, by myself
  - Yes, with coworkers
  - No
- **If you answered yes to the previous question, please briefly explain how: (open-ended):**
- **Has your employer offered any of the following related to AI adoption?**
  - Communicated changes
  - Retraining
  - Legal protections on AI-related layoffs
  - Severance/Transition packages
  - Individual or group counseling
  - Other \_\_\_\_\_
- **Have you taken any of the following steps in response to AI changes at work? (select all that apply)**

- Learned AI tools
- Asked employer for training/support
- Discussed concerns with coworkers
- Raised concerns with a manager/supervisor
- Discussed with my union
- Looked for a different job/career
- No steps taken
- **What has stopped you from taking action around AI changes at work? (select all that apply)**
  - Don't know what steps to take
  - Don't have time/resources
  - Don't think my employer would listen
  - Afraid of being seen as resistant to change
  - Not my responsibility — my employer should handle it
  - Not urgent enough
  - N/A
- **What, if anything, discourages you from challenging how AI/automation is used in your workplace? (open ended)**

### Intervention and Needs

- **Which of the following would be most helpful to you right now? (choose 3 options)**
  - Communicated changes
  - Retraining
  - Legal protections on AI-related layoffs
  - Severance/Transition packages
  - Individual or group counseling
  - Other \_\_\_\_\_
- **Who do you trust most to advocate for workers affected by AI?**
  - Government
  - Nonprofits/Advocacy Groups
  - Unions
  - Managers/Executives
  - HR Departments
  - None of these
  - Unsure