

AI, Jobs, and the Automation Trap: Where Is HCI?

Marios Constantinides
marios.constantinides@cyens.org.cy
CYENS Centre of Excellence
Nicosia, Cyprus

Daniele Quercia
quercia@cantab.net
Nokia Bell Labs
Cambridge, UK
Politecnico di Torino
Turin, Italy

Abstract

As artificial intelligence (AI) continues to reshape the workforce, its current trajectory raises pressing questions about its ultimate purpose. Why does job automation dominate the agenda, even at the expense of human agency and equity? This paper critiques the automation-centric paradigm, arguing that current reward structures, which largely focus on cost reduction, drive the overwhelming emphasis on task replacement in AI patents. Meanwhile, Human-Centered AI (HCAI), which envisions AI as a collaborator augmenting human capabilities and aligning with societal values, remains a fugitive from the mainstream narrative. Despite its promise, HCAI has gone “missing”, with little evidence of its principles translating into patents or real-world impact. To increase impact, actionable interventions are needed to disrupt existing incentive structures within the HCI community. We call for a shift in priorities to support translational research, foster cross-disciplinary collaboration, and promote metrics that reward tangible and real-world impact.

CCS Concepts

• **Human-centered computing** → **HCI theory, concepts and models**; *Collaborative and social computing*.

Keywords

automation, augmentation, artificial intelligence, future of work

ACM Reference Format:

Marios Constantinides and Daniele Quercia. 2025. AI, Jobs, and the Automation Trap: Where Is HCI?. In *CHIWORK '25: Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work (CHIWORK '25)*, June 23–25, 2025, Amsterdam, Netherlands. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3729176.3729191>

1 Introduction

As artificial intelligence (AI) continues to permeate every facet of work, a fundamental question arises: What is its purpose in the workplace? Many AI initiatives today aim to automate tasks of occupations, often with the implicit goal of maximizing efficiency and reducing the need for human involvement [54]. This trajectory

has sparked concerns about widespread job displacement, reduced workforce agency, and exacerbation of existing societal inequalities.

Concerns about automation replacing human labor are not new. The Industrial Revolution offers a compelling historical parallel to today’s AI-driven transformations. Just as 19th-century workers resisted mechanization to preserve their livelihoods, today’s workforce faces growing concerns about AI “taking over” human jobs [19]. Reflecting on that period reveals a key lesson: *when driven solely by efficiency, technological advancement can exacerbate inequality and displace vulnerable workers*. This historical lens sharpens our focus on the present challenge: *will we allow AI to widen socioeconomic divides, or will we intentionally design systems that augment human work and promote inclusion?*

HCAI aims at doing the latter. It emphasizes the design of systems that align with human cognitive strengths, workflows, and societal values, empowering workers rather than marginalizing them [58]. Discussions in special interest groups (SIGs) at recent HCI conferences such as ACM CHI and CSCW [16, 17, 45, 59] have reinforced the critical need to align AI with human cognitive abilities. This alignment is evident in efforts such as adaptive learning platforms that personalize education for diverse learners [18], and AI systems that augment radiologists to achieve greater diagnostic accuracy in collaboration with machine learning models [2].

However, despite decades of HCI work emphasizing participatory and human-centered design, the dominant trajectory of AI design and deployment remains centered on automation (i.e., replacing human tasks) rather than augmentation (i.e., supporting human capabilities). We critically examine this misalignment through three research questions:

RQ1 Does automation dominate AI’s current design ethos, despite long-standing alternatives? We present empirical trends in AI patenting for workplace technologies to map automation-centric innovation trends (Section 2).

RQ2 Are alternatives fostering human augmentation possible? We analyzed case examples of alternatives across HCAI’s three application domains: education, healthcare, and workplace (Section 3). These examples are drawn from the literature based on conceptual relevance and application diversity rather than through a systematic review.

RQ3 How can HCI insights be better translated into real-world AI deployments? We synthesized six actionable recommendations to enhance the translational impact of HCAI research (Section 4). These recommendations emerged from a reflective synthesis of our findings, and were informed by prior work in HCI [11], HCAI [44], and Responsible AI [14].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CHIWORK '25, June 23–25, 2025, Amsterdam, Netherlands

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1384-2/25/06

<https://doi.org/10.1145/3729176.3729191>

2 Does Automation Dominate AI's current Design Ethos?

2.1 Automation Dominates

In two recent papers, we analyzed a dataset of 24,758 AI-related patents filed over the last decade [33, 42]. This analysis was conducted classifying patents according to whether they emphasized automation of job tasks or human involvement. Our goal was not to measure HCI's citation presence per se, but to evaluate the translational trajectory of AI design: do patents encode principles of augmentation or automation?

Our results showed that automation dominates over human augmentation in AI patents because the vast majority of patented AI innovations are aligned with task characteristics that favor replacement rather than collaboration [33]. Specifically, AI patents reinforcing existing technologies and workflows are disproportionately associated with tasks that are predictable and performed individually, aligning with traditional models of automation. Even more disruptive AI patents primarily target mental and unpredictable tasks, yet still favor individual over collaborative work. In contrast, collaborative tasks remain largely untouched by either type of AI, suggesting that augmenting human interaction and cooperation remains a marginal concern in patentable AI development.

To gain a more nuanced and in-depth understanding of how AI patents conceptualize labor, we turned to healthcare as a representative sector. The selection of this domain was based on its rich intersection with both automation and augmentative technologies (e.g., AI is extensively used to automate diagnostic and administrative tasks such as analyzing medical imaging [22, 50], predicting patient outcomes [5], and managing documentation [35]), and because it is a high-stakes domain where the impact of automation versus augmentation has immediate human consequences.

A strong emphasis on automation tasks over augmentation is present in healthcare too (Figure 1). Tasks such as automated analysis of medical imaging (e.g., coronary angiograms and MRI scans), dominate the landscape, as seen with patents focused on technologies such as 'dynamic magnetic resonance imaging' or 'automated vascular analysis'. In contrast, fewer patents target augmentation technologies that assist medical professionals directly (e.g., systems for patient-specific therapy planning or dynamic health record problem lists), which involve supporting rather than replacing human expertise. These findings are in line with recent publications. A report by the American Medical Association found, in fact, that 65% of physicians recognized the utility of AI in reducing administrative burdens, and over half highlighted its potential in diagnostic support [5]. Similarly, a Forbes report noted that AI systems are predominantly applied to automate processes such as disease detection and treatment planning, emphasizing efficiency and scalability over collaboration and augmentation [22]. However, while augmentation-focused applications such as clinical decision support and personalized medicine exist, they represent a smaller portion of AI's integration into healthcare [5].

2.2 Why Does Automation Dominate?

One may now wonder why the dominant focus of AI innovation remains on designing systems that operate autonomously. We offer four complementary explanations.

Because augmenting humans is not easy. Human-centered systems must carefully balance the benefits of AI with concerns around job satisfaction, meaning, and professional identity [26]. For instance, Toner-Rodgers [47] found that pairing researchers with AI systems led to substantial productivity gains: scientists discovered 44% more materials, and patent filings increased by 39%. However, 82% of them also reported feeling less fulfilled in their work. This highlights a paradox: *while AI can boost efficiency and innovation, it may inadvertently diminish the intrinsic rewards of creative and intellectual engagement* [9].

Because long-standing alternatives stem from a field with limited translational impact. HCI research, despite offering robust frameworks for participatory and human-centered design, has historically struggled to influence industry practices. Only 13.4% of publications from SIGCHI-sponsored venues (e.g., CHI, CSCW, UIST, UbiComp) have been cited in AI patents, compared to 25% from fields such as computer vision [11]. This points to a translational gap: a disconnect between academic innovation and its real-world application. While HCI contributions such as direct manipulation [43], heuristic evaluation [39], social translucency [21], and value-sensitive design [24, 40] have shaped theory and practice, they often remain absent from patent-driven innovation ecosystems.

Because automation aligns with dominant worldviews. Historically, men have filed most patents [31] and continue to receive the majority of AI funding [46], embedding particular economic values such as efficiency and control into emerging technologies. As Wajcman argues [52], dominant perspectives in technological development often prioritize performance metrics over relational or affective dimensions of work. As a result, the very notion of "success" in AI-human collaboration is shaped by gendered assumptions about productivity rather than a more inclusive view of human flourishing.

Because automation aligns with short-term economic incentives. The reward structures governing AI innovation typically favor short-term gains: reducing labor costs, scaling operations, and maximizing shareholder value. Yet this economic logic often comes at the expense of long-term societal outcomes. Autor [6, 8] has shown that automation disproportionately targets routine, middle-skill jobs, contributing to job polarization and middle-class erosion. While automation boosts productivity, it narrows pathways to economic mobility. Without intentional investment in technologies that augment human labor, the current trajectory risks exacerbating inequality [7].

3 Any Alternatives Fostering Augmentation?

To address the automation-centric trajectory identified in the previous section, we turn to examples of Human-Centered AI (HCAI) systems that already support human expertise, decision-making,



Figure 1: Jobs impacted by AI in the healthcare sector as visualized in the AI Impact dashboard (<https://social-dynamics.net/aii/>) based on the methodology described by Septiandri et al. [42]. The most impacted jobs are primarily automation-driven, focusing on tasks such as automated vascular analysis and MRI system operations. In contrast, augmentation-driven tasks (e.g., patient-specific therapy planning and dynamic health record management) are less prominent.

and collaboration. Rather than removing human agency, these systems exemplify the shift toward augmentation by helping individuals stay in control, build new skills, and make informed decisions [32, 44]. Building on prior work [15], we analyzed three key domains¹ that are central to the structure of modern labor markets: *education*, *workplace*, and *healthcare*. Education prepares future workers; workplace technologies support those currently in the labor force; and healthcare represents a high-stakes sector undergoing rapid AI-driven change. These domains also reflect many of the 138 mobile and wearable AI use cases recently evaluated against the EU AI Act for their risks and benefits [15]. Across all three, we observed how HCAI principles are being applied to build systems that promote trust, enhance human capability, and serve as viable alternatives to automation-focused deployments.

3.1 Augmenting Learning and Development in Education

Recent studies presented at ACM CHI have showcased innovative educational tools that leverage AI for personalized learning. Although education is not a work setting per se, it plays a foundational role in preparing individuals for a labor market increasingly shaped by AI. Adaptive learning platforms, for example, dynamically adjust content to meet individual student needs that significantly

improve engagement and learning outcomes. For example, Cheng et al. [13] demonstrated the potential of augmented reality (AR) and AI-driven systems to create immersive and culturally relevant learning experiences. Narrative-driven AR applications, coupled with LLMs, could enhance both cognitive and emotional development in children by fostering deep engagement with educational content. Similarly, GPT Coach leverages LLMs to provide tailored tutoring experiences [32]. By incorporating motivational interviewing techniques and leveraging data from wearables, GPT Coach exemplifies how HCAI principles can create personalized and supportive learning environments. Further, GPTeach [36] introduced a novel approach to teacher training by utilizing LLM-powered simulated students to create scalable and low-risk environments for novice educators. Through interactive teaching scenarios, GPTeach enables teachers to refine their instructional strategies, offering immediate feedback and fostering reflective practice. Such tools exemplify the transformative potential of HCAI in education. By personalizing interactions and fostering agency, they empower students and educators alike to achieve deeper learning outcomes and cultivate skills tailored to their individual needs and aspirations. The integration of AR, narrative-based pedagogy, and LLM-driven systems provides a robust pathway for augmenting human capabilities, shifting the focus from replacing educators to augmenting them.

¹The examples in this section are not intended as a systematic review but were selected through a semi-structured search process. We reviewed recent HCI and AI literature with a focus on augmentative design, practical deployment, and cross-sectoral representation. The final selection aimed to balance conceptual diversity, domain relevance, and empirical richness, while illustrating how HCAI principles are applied in practice.

3.2 Augmenting Human Capabilities in the Workplace

AI-driven tools in the workplace are transforming how employees engage with complex tasks, enabling real-time feedback and fostering collaboration. For example, GitHub Copilot and similar code recommendation systems significantly enhance programmer productivity by offering real-time coding suggestions and reducing the cognitive load associated with repetitive tasks [38]. Although AI coding assistants have been shown to expedite task completion for experienced programmers [37], yet their effects on novice programmers remain debated. Some evidence suggests these tools may hinder learning by bypassing foundational understanding [25]. Designing such systems with pedagogical scaffolding and explainability could better support skill development and augmentation in programming tasks. Intelligent systems designed for online meetings further augment collaboration by addressing communication challenges in virtual settings. For example, MeetCues leverages real-time visual and interactive feedback to enhance communication in virtual meetings [4, 60], and KAIROS uses multimodal monitoring (e.g., body language, speech patterns) through wearable devices to capture non-verbal communication cues that otherwise might go unnoticed during virtual meetings. In addition to software engineering and productivity tools, AI is augmenting decision-making in sectors such as finance and law. In finance, AI dashboards assist financial analysts by flagging anomalies and synthesizing data streams to support rather than replacing expert judgment [27]. In the legal domain, AI-powered tools help summarize documents and draft briefs that enable faster iteration while maintaining professional oversight [12]. These systems exemplify how AI can offload repetitive tasks while preserving human responsibility and ethical accountability [3].

3.3 Augmenting Decision-Making in Healthcare

AI systems such as diagnostic assistants analyze patient data to provide clinicians with actionable insights, helping to improve decision-making in high-stakes environments. For example, it has been shown that adapted LLMs can outperform medical experts in clinical text summarization tasks, including summarizing radiology reports, patient questions, progress notes, and doctor–patient dialogue [50]. As LLM-generated summaries were rated as equivalent or superior to those created by medical experts in terms of completeness, correctness, and conciseness, by integrating such models into clinical workflows could alleviate the documentation burden and allow clinicians to allocate more time to patient care. However, as Autor [6] highlights, the effectiveness of such tools depends on the ability of healthcare professionals to understand and utilize them appropriately. For example, a recent study by Nikhil Agarwal et al. demonstrated that while AI diagnostic tools performed at least as accurately as nearly two-thirds of radiologists in an experimental setting, they did not improve the quality of diagnoses. The key issue was that radiologists often misinterpreted or misused the AI's predictions. Confident AI predictions were frequently overridden by clinicians, and when AI expressed uncertainty, doctors often relied on these predictions even when their own assessments were better [2]. Research evidence also suggests that AI tools can help reduce burnout among healthcare professionals by automating

routine tasks and alleviating cognitive load [35]. Centered on augmentation rather than replacement, this example illustrates how HCAI principles can guide the development of AI systems that enhance both patient care and healthcare worker well-being.

4 How Can HCI Translate into Real-world AI Deployments?

To translate the principles of HCAI into real-world impact, we must align ethical imperatives with institutional incentives. In today's political and business climate, where efficiency, scalability, and profitability often dominate, the success of augmentative AI systems depends not only on moral alignment, but also on strategic advantage. Based on our review of patent trends, case studies, and the broader HCI literature, we propose six recommendations to increase the translational impact (Table 1). These recommendations are not prescriptive blueprints, but synthesized themes that emerged from cross-domain analysis and our own practitioner experience in HCI research and design. We present them here as entry points for discussion and iteration.

4.1 Taking a “Responsible AI by Design” Approach

“Responsible AI by Design” involves building AI systems responsibly from the very beginning, starting in the design phase, to prevent issues later in development. Developers and industry leaders should prioritize safety from the outset, taking inspiration from how self-driving car companies integrate safety measures at every stage of their design process. This approach ensures that AI aligns with human needs, integrates seamlessly into daily life, minimizes risks, and remains productive and helpful. By adopting a structured approach [14], developers can establish clear Responsible AI (RAI) guidelines. These guidelines, grounded in regulations such as the EU AI Act, help ensure that safety and fairness are embedded throughout the development process [56]. A practical way to implement this is by incorporating RAI guidelines into AI workflows using interactive tools and visual dashboards. For example, digital prompts or cards that highlight ethical considerations at each stage of development can promote adherence to core principles such as fairness, transparency, and safety [20, 49]. While ethical guidelines such as fairness and transparency are essential, they must also be framed in terms of operational and reputational benefits. Research shows that responsible design mitigates costly downstream failures, reduces compliance burdens, and strengthens public trust [14, 20, 34, 49]; factors that increasingly affect brand perception and market share. As the regulatory landscapes evolve (e.g., EU AI Act), adopting these principles early can serve as a strategic safeguard against reputational or legal harm.

4.2 Fostering Academia-Industry Collaboration for Workforce Innovation

To ensure HCAI research addresses real-world job needs, researchers should prioritize partnerships with industries where AI is already transforming work (e.g., manufacturing, healthcare, and service sectors) [8, 42]. Joint initiatives, such as workforce retraining programs or augmented workplace studies, can inform research agendas and

Table 1: Summary of recommendations for increasing HCAI's translational impact.

Recommendation	Key Actions	Expected Impacts
Responsible AI by Design	- Incorporate ethical guidelines (e.g., fairness, transparency) into AI workflows. - Use tools like interactive dashboards to monitor adherence to principles.	- Ensures AI systems align with regulatory frameworks (e.g., EU AI Act). - Reduces risks of bias and enhances public trust.
Fostering Academia-Industry Collaboration	- Develop partnerships with industries such as healthcare and manufacturing. - Establish research labs embedded in companies to align with real-world needs.	- Facilitates faster adoption of HCAI innovations. - Informs research agendas based on practical workflows.
Developing Scalable Prototypes for Workforce Applications	- Create job-specific AI prototypes addressing key industry needs (e.g., healthcare diagnostics). - Make prototypes available as open-source tools to encourage broader use.	- Provides clear examples of augmentative AI in action. - Encourages collaboration and faster scaling across industries.
Targeting Job Preservation and Skill Development	- Design AI systems to complement human labor in automation-prone sectors. - Develop adaptive learning platforms for reskilling and upskilling workers.	- Mitigates job displacement by preparing workers for AI-augmented roles. - Addresses labor shortages in caregiving, education, and logistics.
Embedding Feedback Loops to Address Worker Concerns	- Use participatory design to include workers in system development. - Implement contestability mechanisms to enable workers to challenge AI decisions.	- Ensures systems align with worker needs and reduce physical or cognitive strain. - Increases trust and satisfaction with AI tools.
Shaping Policy for Augmentative AI	- Collaborate with government agencies to prioritize augmentation in AI policies. - Advocate for regulations requiring participatory design and human oversight.	- Aligns AI deployment with societal values like fairness and inclusivity. - Encourages policies that reduce job polarization.

facilitate faster adoption of innovations. Establishing research labs embedded within companies could help HCAI researchers gain deeper insights into job workflows, enabling them to design AI systems that genuinely augment human labor. While establishing in-house research labs may appear costly, the long-term return on investment can be substantial. Embedded collaboration helps companies avoid deployment failures by aligning AI systems with actual work practices, reducing workforce resistance, and enhancing employee retention. Similar to how usability testing reduces product returns and support costs, HCI-informed augmentation strategies can reduce turnover and burnout while improving system adoption. As Burke et al. [10] show, systems designed with social and well-being outcomes in mind can yield broader engagement benefits. Furthermore, evidence from organizational research suggests that collaborative and participatory design processes improve trust, which is vital for AI adoption in high-stakes domains [48].

Cross-disciplinary collaborations between HCI researchers, social scientists, economists, and policymakers are also essential [53]. For example, partnerships with labor economists can provide evidence of the economic benefits of augmentation, while collaborations with policymakers can ensure these insights influence regulatory frameworks. Future work in this area could focus on making these guidelines easier to use in real-world AI applications, developing tools that help developers create AI systems that enhance rather than replace human jobs, and studying how these

guidelines impact employment opportunities and workforce dynamics across different industries. Additionally, as academic research often emphasizes novelty over practical impact, it is important to develop new reward structures that incentivize researchers to perform translational work [11, 41].

4.3 Developing Scalable Prototypes for Workforce Applications

HCAI researchers should focus on developing scalable and job-specific AI prototypes that offer clear benefits by augmenting workers' capabilities. Ideally, these prototypes should be made available as open-source tools to encourage broader adoption and collaboration. For example, ExploreGen [29] demonstrates how generative AI can support researchers in quickly creating prototypes while simultaneously identifying potential risks and benefits, ensuring alignment with human-centered goals and responsible design principles. Such prototypes can include AI tools that assist mid-skill workers, like diagnostic support systems for medical technicians, enabling them to make better decisions. Finally, AI systems designed to reduce cognitive load can streamline complex tasks for professionals in high-stakes environments, such as air traffic controllers or radiologists, enhancing both accuracy and efficiency [15, 29].

4.4 Empowering Workers Through Job Preservation and Skill Development

AI systems that disproportionately automate routine tasks have contributed to the erosion of middle-skill jobs, leading to job polarization and economic inequality [1, 28, 42]. In roles vulnerable to automation, AI systems should complement human labor rather than replacing it. However, we acknowledge that in certain high-risk environments (e.g., manufacturing or chemical processing) automation can increase worker safety by removing people from hazardous conditions. For example, replacing human labor on factory floors with AI-driven machines can prevent accidents involving heavy machinery or toxic exposure. Yet even in these contexts, hybrid systems that maintain human oversight or offer augmentation through robotics-assisted safety mechanisms can strike a more balanced approach between efficiency and worker agency. Another way would be to integrate skill development initiatives into the design and deployment of these systems, ensuring workers are equipped to transition into higher-value tasks within their existing roles. For example, AI-driven training simulators and adaptive learning tools can help workers acquire skills tailored to emerging demands in industries such as manufacturing and healthcare [30]. These systems can bridge the gap between traditional education and augmented workflows, preparing workers to thrive in AI-augmented roles. Additionally, by collaborating with educational institutions, HCAI researchers can promote AI literacy and integrate vocational training programs that address the specific needs of AI-impacted sectors. Moreover, addressing labor shortages in critical areas such as caregiving and education requires designing AI systems that enhance efficiency without compromising the essential human qualities of empathy and creativity. For example, augmentation tools that assist caregivers with administrative tasks can enable them to focus on delivering personalized care. While previous literature has suggested reallocating low-skilled workers to tasks requiring creative and social intelligence [23], we argue that (re)skilling and upskilling efforts should be approached cautiously, ensuring that new AI systems align with workers' capabilities and societal needs. For example, while a truck driver may not transition to an electronics technician, they might excel as a logistics coordinator or sales professional. Similarly, a logistics coordinator transitioning from truck driving can benefit from AI-driven scheduling systems that enhance decision-making without requiring extensive retraining. At the same time, opportunities for (re)skilling and upskilling can also foster joy, self-efficacy, and personal development; factors that are critical for long-term worker engagement and well-being. Future work should explore the diverse aspirations and capacities of workers across industries to design augmentation-focused interventions that resonate with their needs.

4.5 Embedding Feedback Loops to Address Worker Concerns

To align HCAI research with job-centric challenges, embedding workers and managers into the feedback loops of AI development processes is essential. This approach ensures that AI systems not only meet technical requirements but also address the practical needs and concerns of those who interact with them daily. A key component of such integration is the concept of AI contestability,

which emphasizes designing systems that allow decisions and processes to be challenged or adapted by those directly or indirectly affected [3]. For example, involving warehouse workers in the participatory design of robotics-assisted picking systems ensures that these technologies align with ergonomic needs, reducing physical strain and improving productivity. Similarly, engaging healthcare workers in designing AI-assisted diagnostics tools fosters systems that fit seamlessly into clinical workflows, minimizing disruption and improving user satisfaction [51]. Such participatory methods create a collaborative feedback loop where stakeholders can contest and refine system decisions, leading to tools that are both effective and ethically aligned with human needs. Another way of allowing workers to understand AI's impact is through visualization tools that highlight disparities caused by automation (e.g., job displacement rates across industries or demographic groups).

4.6 Shaping Policy for Augmentative AI

HCAI researchers can significantly increase their impact by shaping policy frameworks that guide the deployment of AI in the workforce. Recent insights from special interest groups organized at ACM CHI and CSCW have highlighted how regulatory policies such as the EU AI Act can support augmentative AI systems by mandating transparency, user empowerment, and the ability for individuals to contest AI decisions [16, 17, 45]. These policies align with Human-Centered AI principles, emphasizing the need for AI systems to augment human labor rather than replace it. However, influencing policy requires HCI to consider policy engagement not as peripheral but as central to its intellectual and practical concerns. As recent work suggests, HCI can expand its role by positioning system-people-policy interactions at the core of its research, practice, and education [55]. This approach allows HCI to blend system, human, and policy expertise seamlessly, creating a cohesive framework for addressing the societal implications of AI. Such a shift enables the HCAI community to move beyond isolated policy engagements and work collectively to influence outcomes that prioritize human augmentation. Collaborating with government agencies, labor unions, and industry stakeholders, HCAI researchers can advocate for policies that encourage participatory design and ensure AI systems enhance worker well-being. For example, promoting mandates that require human oversight in AI-driven decision-making processes can foster trust and usability while mitigating risks associated with automation [3, 57]. Furthermore, HCI researchers can provide empirical evidence on the benefits of augmentative AI, drawing on case studies where AI has successfully enhanced human capabilities in healthcare, education, and other critical sectors.

5 Conclusion

Our analysis demonstrates a systemic disconnect between HCI's augmentation-oriented ethos and AI's current automation-first trajectory. While HCAI has emerged as a compelling framework, the path from design principle to industrial deployment remains unclear. Part of the challenge lies in incentive structures: patents, funding models, and commercial benchmarks often reward cost-reduction and task elimination rather than human empowerment or resilience.

Moreover, the line between automation and augmentation is blurry. In some domains, automation may in fact enable augmentation by, for example, reducing cognitive load or freeing time for complex tasks. However, this requires intentional design. Without proactive interventions, automation risks degrading work quality or removing workers from critical feedback loops, thereby increasing system fragility.

The recommendations outlined in this paper, ranging from embedding feedback loops to advocating for augmentative AI policies, highlight actionable pathways for increasing HCAI's translational impact. We proposed strategies that emphasize collaboration with diverse stakeholders, including workers, educational institutions, and policymakers, to ensure that AI aligns with human values and societal needs. By doing so, HCAI will be positioned not only as a moral imperative, but as a competitive advantage.

AI USAGE DISCLOSURE STATEMENT

In the writing and revision stages, we used a generative AI tool (i.e., OpenAI's ChatGPT-4o) to support activities such as refining the clarity of our arguments and enhancing overall readability. The tool was used as a writing aid and did not generate original scientific content or research findings.

References

- [1] Daron Acemoglu and Pascual Restrepo. 2022. Tasks, automation, and the rise in US wage inequality. *Econometrica* 90, 5 (2022), 1973–2016. <https://doi.org/10.3982/ECTA19815>
- [2] Nikhil Agarwal, Alex Moehring, Pranav Rajpurkar, and Tobias Salz. 2023. *Combining human expertise with artificial intelligence: Experimental evidence from radiology*. Technical Report. National Bureau of Economic Research.
- [3] Kars Alfrink, Ianus Keller, Gerd Kortuem, and Neelke Doorn. 2023. Contestable AI by design: towards a framework. *Minds and Machines* 33, 4 (2023), 613–639. <https://doi.org/10.1007/s11023-022-09611-z>
- [4] Bon Adriel Aseniero, Marios Constantinides, Sagar Joglekar, Ke Zhou, and Daniele Quercia. 2020. MeetCues: Supporting online meetings experience. In *IEEE Visualization Conference (VIS)*. IEEE, 236–240. <https://doi.org/10.1109/VIS47514.2020.00054>
- [5] American Medical Association et al. 2023. AMA Augmented Intelligence Research: Physician Sentiments around the Use of AI in Health Care: Motivations, Opportunities, Risks, and Use Cases.
- [6] David Autor. 2024. *AI Could Actually Help Rebuild The Middle Class*. <https://www.noemamag.com/how-ai-could-help-rebuild-the-middle-class/>
- [7] David Autor. 2024. *Applying AI to rebuild middle class jobs*. Technical Report. National Bureau of Economic Research.
- [8] David Autor, Caroline Chin, Anna Salomons, and Bryan Seegmiller. 2024. New frontiers: The origins and content of new work, 1940–2018. *The Quarterly Journal of Economics* (2024). <https://doi.org/10.1093/qje/qjae008>
- [9] Lisanne Bainbridge. 1983. Ironies of automation. In *Analysis, design and evaluation of man-machine systems*. Elsevier, 129–135.
- [10] Moira Burke, Cameron Marlow, and Thomas Lento. 2010. Social network activity and social well-being. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1909–1912.
- [11] Hancheng Cao, Yujie Lu, Yuting Deng, Daniel Mcfarland, and Michael S Bernstein. 2023. Breaking out of the ivory tower: A large-scale analysis of patent citations to HCI research. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. 1–24. <https://doi.org/10.1145/3544548.3581108>
- [12] Longbing Cao. 2022. AI in finance: challenges, techniques, and opportunities. *ACM Computing Surveys (CSUR)* 55, 3 (2022), 1–38.
- [13] Alan Y Cheng, Meng Guo, Melissa Ran, Arpit Ranasaria, Arjun Sharma, Anthony Xie, Khuyen N Le, Bala Vinaitirthan, Shih Luan, David Thomas Henry Wright, et al. 2024. Scientific and Fantastical: Creating Immersive, Culturally Relevant Learning Experiences with Augmented Reality and Large Language Models. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. 1–23. <https://doi.org/10.1145/3613904.3642041>
- [14] Marios Constantinides, Edyta Bogucka, Daniele Quercia, Susanna Kallio, and Mohammad Tahaei. 2024. RAI Guidelines: Method for Generating Responsible AI Guidelines Grounded in Regulations and Usable by (Non-) Technical Roles. *Proceedings of the ACM on Human-Computer Interaction (CSCW)* 8, CSCW2 (2024), 1–28. <https://doi.org/10.1145/3686927>
- [15] Marios Constantinides, Edyta Paulina Bogucka, Sanja Scepanovic, and Daniele Quercia. 2024. Good Intentions, Risky Inventions: A Method for Assessing the Risks and Benefits of AI in Mobile and Wearable Uses. *Proceedings of the ACM on Human-Computer Interaction (MobileHCI)* 8, MHCI (2024), 1–28. <https://doi.org/10.1145/3676507>
- [16] Marios Constantinides, Mohammad Tahaei, Daniele Quercia, Simone Stumpf, Michael Madaio, Sean Kennedy, Lauren Wilcox, Jessica Vitak, Henriette Cramer, Edyta Paulina Bogucka, et al. 2024. Implications of Regulations on the Use of AI and Generative AI for Human-Centered Responsible Artificial Intelligence. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems (CHI)*. 1–4. <https://doi.org/10.1145/3613905.3643979>
- [17] Wesley Hanwen Deng, Glen Berman, Harmanpreet Kaur, Reva Schwartz, Munat A Ibrahim, Motahhare Eslami, Kenneth Holstein, and Michael Madaio. 2024. Collaboratively Designing and Evaluating Responsible AI Interventions. In *Companion Publication of the ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*. 658–662. <https://doi.org/10.1145/3678884.3687136>
- [18] Srijani Dutta, Sakshi Ranjan, Sushruta Mishra, Vandana Sharma, Pradeep Hewage, and Celestine Iwendu. 2024. Enhancing educational adaptability: A review and analysis of AI-driven adaptive learning platforms. In *International Conference on Innovative Practices in Technology and Management (ICIPTM)*. IEEE, 1–5. <https://doi.org/10.1109/ICIPTM59628.2024.10563448>
- [19] Matt Egan. 2024. *AI is replacing human tasks faster than you think*. CNN. <https://edition.cnn.com/2024/06/20/business/ai-jobs-workers-replacing/index.html>
- [20] Salma Elsayed-Ali, Sara E. Berger, Vagner F. De Santana, and Juana Catalina Becerra Sandoval. 2023. Responsible & Inclusive Cards: An Online Card Tool to Promote Critical Reflection in Technology Industry Work Practices. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. 1–14. <https://doi.org/10.1145/3544548.3580771>
- [21] Thomas Erickson and Wendy A Kellogg. 2000. Social translucence: an approach to designing systems that support social processes. *ACM transactions on computer-human interaction (TOCHI)* 7, 1 (2000), 59–83.
- [22] Forbes. 2024. 19 Major Impacts AI Is Having On The Healthcare Industry. <https://www.forbes.com/councils/forbesbusinesscouncil/2024/09/19/19-major-impacts-ai-is-having-on-the-healthcare-industry/>
- [23] Carl Benedikt Frey and Michael A Osborne. 2017. The future of employment: How susceptible are jobs to computerisation? *Technological forecasting and social change* 114 (2017), 254–280. <https://doi.org/10.1016/j.techfore.2016.08.019>
- [24] Batya Friedman. 1996. Value-sensitive design. *Interactions* 3, 6 (1996), 16–23.
- [25] Nicholas Gardella, Raymond Pettit, and Sara L Riggs. 2024. Performance, Workload, Emotion, and Self-Efficacy of Novice Programmers Using AI Code Generation. In *Proceedings of the 2024 on Innovation and Technology in Computer Science Education V. 1*. 290–296.
- [26] Kuntal Ghosh and Shadan Sadeghian. 2024. The Impact of AI on Perceived Job Decency and Meaningfulness: A Case Study. *arXiv preprint arXiv:2406.14273* (2024).
- [27] Peter Gomber, Robert J Kauffman, Chris Parker, and Bruce W Weber. 2018. On the fintech revolution: Interpreting the forces of innovation, disruption, and transformation in financial services. *Journal of management information systems* 35, 1 (2018), 220–265.
- [28] Maarten Goos, Alan Manning, and Anna Salomons. 2009. Job Polarization in Europe. *American Economic Review* 99, 2 (April 2009), 58–63. <https://doi.org/10.1257/aer.99.2.58>
- [29] Viviane Herdel, Sanja Šćepanović, Edyta Bogucka, and Daniele Quercia. 2024. ExploreGen: Large Language Models for Envisioning the Uses and Risks of AI Technologies. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES)*, Vol. 7. 584–596. <https://doi.org/10.1609/aies.v7i1.31660>
- [30] Shanti Jagannathan, Sungsup Ra, and Rupert Maclean. 2019. Dominant recent trends impacting on jobs and labor markets—An Overview. *International Journal of Training Research* 17, sup1 (2019), 1–11. <https://doi.org/10.1080/14480220.2019.1641292>
- [31] Kyle Jensen, Balázs Kovács, and Olav Sorenson. 2018. Gender differences in obtaining and maintaining patent rights. *Nature biotechnology* 36, 4 (2018), 307–309.
- [32] Matthew Jörke, Shardul Sapkota, Lyndsea Warkenthien, Niklas Vainio, Paul Schmiedmayer, Emma Brunskill, and James Landay. 2024. Supporting Physical Activity Behavior Change with LLM-Based Conversational Agents. *arXiv preprint arXiv:2405.06061* (2024).
- [33] Munjung Kim, Marios Constantinides, Sanja Scepanovic, Yong-Yeol Ahn, and Daniele Quercia. 2025. The Potential Impact of Disruptive AI Innovations on U.S. Occupations. *Arxiv* (2025).
- [34] Q Vera Liao, Daniel Gruen, and Sarah Miller. 2020. Questioning the AI: informing design practices for explainable AI user experiences. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–15.
- [35] Hanyang Liu, Sunny S Lou, Benjamin C Warner, Derek R Harford, Thomas Kannampallil, and Chenyang Lu. 2022. Hipal: A deep framework for physician burnout prediction using activity logs in electronic health records. In *Proceedings*

- of the ACM Conference on Knowledge Discovery and Data Mining (KDD). 3377–3387. <https://doi.org/10.1145/3534678.3539056>
- [36] Julia M Markel, Steven G Opferman, James A Landay, and Chris Piech. 2023. Gpteach: Interactive ta training with gpt-based students. In *Proceedings of the ACM Conference on Learning@ Scale*. 226–236. <https://doi.org/10.1145/3573051.3593393>
- [37] Ran Mo, Dongyu Wang, Wenjing Zhan, Yingjie Jiang, Yepeng Wang, Yuqi Zhao, Zengyang Li, and Yutao Ma. [n. d.]. Assessing and Analyzing the Correctness of GitHub Copilot's Code Suggestions. *ACM Transactions on Software Engineering and Methodology* ([n. d.]).
- [38] Hussein Mozannar, Gagan Bansal, Adam Fourney, and Eric Horvitz. 2024. Reading between the lines: Modeling user behavior and costs in AI-assisted programming. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. 1–16. <https://doi.org/10.1145/3613904.3641936>
- [39] Jakob Nielsen and Rolf Molich. 1990. Heuristic evaluation of user interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 249–256.
- [40] Malak Sadek, Marios Constantinides, Daniele Quercia, and Celine Mougenot. 2024. Guidelines for Integrating Value Sensitive Design in Responsible AI Toolkits. In *Proceedings of the 2024 ACM Conference on Human Factors in Computing Systems (CHI)*. Article 472, 20 pages. <https://doi.org/10.1145/3613904.3642810>
- [41] Ali Akbar Septiandri, Marios Constantinides, and Daniele Quercia. 2024. The Impact of Responsible AI Research on Innovation and Development. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES)*, Vol. 7. 1329–1342. <https://doi.org/10.1609/aies.v7i1.31727>
- [42] Ali Akbar Septiandri, Marios Constantinides, and Daniele Quercia. 2024. The potential impact of AI innovations on US occupations. *PNAS Nexus* 3, 9 (2024), pgae320. <https://doi.org/10.1093/pnasnexus/pgae320>
- [43] Ben Shneiderman. 1983. Direct manipulation: A step beyond programming languages. *Computer* 16, 08 (1983), 57–69.
- [44] Ben Shneiderman. 2020. Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human–Computer Interaction* 36, 6 (2020), 495–504. <https://doi.org/10.1080/10447318.2020.1741118>
- [45] Mohammad Tahaei, Marios Constantinides, Daniele Quercia, Sean Kennedy, Michael Muller, Simone Stumpf, Q Vera Liao, Ricardo Baeza-Yates, Lora Aroyo, Jess Holbrook, et al. 2023. Human-Centered Responsible Artificial Intelligence: Current & Future Trends. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems (CHI)*. 1–4. <https://doi.org/10.1145/3544549.3583178>
- [46] The Alan Turing Institute. 2024. Rebalancing Innovation: Women, AI and Venture Capital in the UK. <https://www.turing.ac.uk/news/publications/rebalancing-innovation>
- [47] Aidan Toner-Rodgers. 2024. Artificial Intelligence, Scientific Discovery, and Product Innovation. (2024).
- [48] Michelle Vaccaro, Abdullah Almaatouq, and Thomas Malone. 2024. When combinations of humans and AI are useful: A systematic review and meta-analysis. *Nature Human Behaviour* (2024), 1–11.
- [49] Ville Vakkuri, Kai-Kristian Kemell, Marianna Jantunen, Erika Halme, and Pekka Abrahamsson. 2021. ECCOLA — A Method for Implementing Ethically Aligned AI Systems. *Journal of Systems and Software* 182 (2021), 111067. <https://doi.org/10.1016/j.jss.2021.111067>
- [50] Dave Van Veen, Cara Van Uden, Louis Blankemeier, Jean-Benoit Delbrouck, Asad Aali, Christian Bluethgen, Anuj Pareek, Malgorzata Polacin, Eduardo Pontes Reis, Anna Seehofnerová, et al. 2024. Adapted large language models can outperform medical experts in clinical text summarization. *Nature Medicine* 30, 4 (2024), 1134–1142. <https://doi.org/10.1038/s41591-024-02855-5>
- [51] Himanshu Verma, Jakub Mlynar, Roger Schaer, Julien Reichenbach, Mario Jreige, John Prior, Florian Evéquoz, and Adrien Depeursinge. 2023. Rethinking the role of AI with physicians in oncology: revealing perspectives from clinical and research workflows. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. 1–19. <https://doi.org/10.1145/3544548.3581506>
- [52] Judy Wajcman. 2007. From women and technology to gendered technoscience. *Information, Community and Society* 10, 3 (2007), 287–298.
- [53] Hanna Wallach, Meera Desai, A Feder Cooper, Angelina Wang, Chad Atalla, Solon Barocas, Su Lin Blodgett, Alexandra Chouldechova, Emily Corvi, P Alex Dow, et al. 2025. Position: Evaluating Generative AI Systems is a Social Science Measurement Challenge. *arXiv preprint arXiv:2502.00561* (2025).
- [54] World Economic Forum. 2025. *The Future of Jobs Report 2025*. Technical Report. World Economic Forum. <https://www.weforum.org/publications/the-future-of-jobs-report-2025/>
- [55] Qian Yang, Richmond Y Wong, Steven Jackson, Sabine Junginger, Margaret D Hagan, Thomas Gilbert, and John Zimmerman. 2024. The Future of HCI-Policy Collaboration. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. 1–15. <https://doi.org/10.1145/3613904.3642771>
- [56] Sofia Yfantiidou, Marios Constantinides, Dimitris Spathis, Athena Vakali, Daniele Quercia, and Fahim Kawsar. 2023. Beyond accuracy: a critical review of fairness in machine learning for mobile and wearable computing. *arXiv preprint arXiv:2303.15585* (2023).
- [57] Mireia Yurrita, Tim Draws, Agathe Balayn, Dave Murray-Rust, Nava Tintarev, and Alessandro Bozzon. 2023. Disentangling fairness perceptions in algorithmic decision-making: the effects of explanations, human oversight, and contestability. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. 1–21. <https://doi.org/10.1145/3544548.3581161>
- [58] Daniel Zhang, Saurabh Mishra, Erik Brynjolfsson, John Etchemendy, Deep Ganguli, Barbara Grosz, Terah Lyons, James Manyika, Juan Carlos Nieves, Michael Sellitto, Yoav Shoham, Jack Clark, and Raymond Perrault. 2021. The AI Index 2021 Annual Report.
- [59] Qingxiao Zheng, Xi Lu, Qiao Jin, Jitesh Jain, Hedda Meadan-Kaplansky, Humphrey Shi, Jinjun Xiong, and Yun Huang. 2024. Towards Responsible Use of Large Multi-modal AI to Analyze Human Social Behaviors. In *Companion Publication of the 2024 Conference on Computer-Supported Cooperative Work and Social Computing*. 663–665. <https://doi.org/10.1145/3678884.3687137>
- [60] Ke Zhou, Marios Constantinides, Sagar Joglekar, and Daniele Quercia. 2022. Predicting meeting success with nuanced emotions. *IEEE Pervasive Computing* 21, 2 (2022), 51–59. <https://doi.org/doi={10.1109/MPRV.2022.3145047}>