

# Agentic Automation Experiences—Rethinking the Interaction of Humans and AI Agents

Philipp Spitzer  
Karlsruhe Service Research Institute  
Karlsruhe Institute of Technology  
Karlsruhe, Germany  
philipp.spitzer@kit.edu

Matthias Baldauf  
Eastern Switzerland University of  
Applied Sciences  
St.Gallen, Switzerland  
matthias.baldauf@ost.ch

Philippe Palanque  
ICS-IRIT  
University Paul Sabatier - Toulouse III  
Toulouse, France  
palanque@irit.fr

Virpi Roto  
School of Arts, Design and  
Architecture  
Aalto University  
Espoo, Finland  
virpi.roto@aalto.fi

Katelyn Morrison  
Human-Computer Interaction  
Institute, School of Computer Science,  
DIG Lab  
Carnegie Mellon University  
Pittsburgh, Pennsylvania, USA  
kcmorris@andrew.cmu.edu

Garoa Gomez-Beldarrain  
Industrial Design Engineering faculty  
TU Delft  
Delft, Netherlands  
g.gomezbeldarrain@tudelft.nl

Monika Westphal  
IE University  
Madrid, Spain  
monika.westphal@ie.edu

Joshua Holstein  
Karlsruhe Institute of Technology  
Karlsruhe, Germany  
joshua.holstein@kit.edu

## Abstract

Recent advances in Artificial Intelligence (AI) have enabled agentic AI systems that coordinate multiple, specialized agents behind unified interfaces. These systems can independently initiate actions and solve complex problems. In traditional automation systems within organizations, workers maintained clear oversight—they could see which system handled each task and trace outcomes to specific processes. The integration of agentic AI, however, obscures this relationship and makes it more difficult for humans to identify which agent is responsible for a given outcome. This creates novel research challenges in the field of “Automation Experience”, particularly in terms of transparency, human agency, and long-term human-AI collaboration dynamics. This workshop focuses on these three critical research dimensions. First, multi-agent transparency and attribution explore how humans understand decision-making when responsibility is shared across multiple coordinating agents. Second, human agency examines how workers can keep control when collaborating with proactive AI systems that act on their own. Third, long-term temporal evolution looks at human skills change over time, including how skills are maintained and how dependencies form. Through real-life organizational cases, presentations, and collaborative activities, workshop participants will advance their understanding of human experience with agentic AI and establish a research agenda for organizational contexts.

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).

*CHI EA '26, Barcelona, Spain*

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2281-3/26/04

<https://doi.org/10.1145/3772363.3778732>

## CCS Concepts

• **Human-centered computing** → **Human computer interaction (HCI)**.

## Keywords

Automation, User Experience, Agentic AI, Transparency, Agency

## ACM Reference Format:

Philipp Spitzer, Matthias Baldauf, Philippe Palanque, Virpi Roto, Katelyn Morrison, Garoa Gomez-Beldarrain, Monika Westphal, and Joshua Holstein. 2026. Agentic Automation Experiences—Rethinking the Interaction of Humans and AI Agents. In *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3772363.3778732>

## 1 Motivation

### 1.1 Introduction

Recent advances in Artificial Intelligence (AI) have shifted its capabilities from performing narrow, well-defined tasks to demonstrating enhanced reasoning, creativity, and adaptability. These advances have enabled agentic AI systems that can initiate actions, reason through complex problems independently, and coordinate multiple specialized agents to achieve broader goals [2]. Unlike traditional automation that requires human initiation and operates within predefined boundaries, agentic AI can proactively engage with users [24], adjust its strategies based on changing situations [7], and orchestrate multiple specialized agents behind unified interfaces [11, 14, 33], thereby fundamentally reshaping how humans experience and engage with AI systems in their daily work [26].

This technological evolution presents fundamental challenges for automation experience research. Traditionally, humans are expected to maintain direct oversight of automated systems [3]. System behavior is assumed to be relatively predictable, and workers remain essential for various activities, including maintenance, command, or repair [9, 16, 21]. Thus, workers have to understand the relationship between system inputs and resulting outputs. Agentic AI disrupts these assumptions: humans may not know when or why specific actions occur, system behavior evolves through learning and adaptation, and multi-agent architectures obscure decision-making processes. Within organizations, these challenges intensify as AI operates across departmental boundaries, makes autonomous decisions that affect multiple stakeholders, and coordinates with external systems [20]. The automation experience, thus, becomes deeply embedded within organizational structures, shaped by real-world operational constraints and socio-technical interdependencies, existing hierarchies and workflows, and the ongoing processes through which workers learn to work alongside these systems [15, 19, 29, 34].

Understanding these dynamics requires examining three critical dimensions through situated, real-world organizational cases—specifically workplace contexts where agentic AI systems integrate into existing hierarchies, workflows, and role structures, affecting how workers perform tasks and maintain expertise: **“Enabling Multi-Agent Transparency and Attribution”** explores how humans understand systems where multiple specialized agents coordinate behind unified interfaces, making it difficult to discern responsibility for specific outcomes. **“Balancing Human Agency in Agentic AI”** investigates how workers experience their sense of agency when collaborating with systems that independently initiate actions and make decisions across organizational boundaries. **“Sustaining Human Skills in Long-Term Human-AI Interaction”** examines how relationships between humans, agentic AI systems, and organizational structures develop and transform as technology and work practices co-evolve.

## 1.2 Workshop Goals

The workshop addresses these three challenges through an interdisciplinary approach, building both theoretical understanding and practical design strategies for exploring agentic automation experiences. The workshop objectives are to:

- Examine challenges and opportunities in interaction with multi-agent systems,
- Develop theoretical frameworks for understanding agency and temporal evolutions in organizational agentic AI systems,
- Identify promising future research topics in the form of project ideas and a research agenda,
- Expand and strengthen a multi-disciplinary network of automation experience researchers.

## 2 Challenges

Agentic AI systems pose new challenges for the “Automation Experience” field, which we categorize in the following three categories, namely, enabling multi-agent transparency and attribution (Section 2.1), balancing human agency in agentic AI (Section 2.2), and

sustaining human skills in long-term human-AI interaction (Section 2.3).

### 2.1 Enabling Multi-Agent Transparency and Attribution

Understanding how systems work and being able to attribute outcomes to specific processes and systems has been fundamental to effective human-AI interaction. This transparency enables humans to build appropriate reliance [31], identify when systems fail, debug problematic behavior, and maintain accountability for system outcomes [28]. In real-world contexts, the ability to make decisions and behaviors of specific system components interpretable becomes critical for compliance, quality control, and stakeholder communication [27].

Multi-agent agentic AI systems fundamentally transform existing transparency challenges by distributing decision-making across multiple specialized agents that coordinate behind unified user interfaces [8]. Humans interact with what appears to be a single system, but their requests may be handled by multiple agents, including knowledge retrieval agents, reasoning agents, and execution agents working in dynamic coordination. When outcomes emerge from these interactions, humans cannot determine which agents contributed to the result and how their coordination influences the final output, or which agents should be examined if problems arise. This creates novel attribution challenges where the very notion of identifying “the agent responsible” becomes increasingly difficult.

These multi-agent attribution problems present specific challenges that differ from general transparency issues: organizations cannot effectively debug multi-agent system behavior, cannot assign development requests for problematic outcomes, and cannot provide stakeholders with explanations of how decisions were reached when multiple invisible agents contributed to organizational actions. Research questions include:

- What attribution mechanisms can help humans understand which agents contribute to specific outcomes in multi-agent systems?
- How much transparency into multi-agent coordination do humans need for effective collaboration, and when does additional transparency become counterproductive?
- How can interfaces be designed to provide transparency into multi-agent coordination without overwhelming humans?
- What debugging and intervention approaches work effectively when problems emerge from distributed agent interactions?

### 2.2 Balancing Human Agency in Agentic AI

Human agency, the capacity to exercise control over one’s work processes, represents a fundamental psychological need that impacts job satisfaction, professional identity, and organizational effectiveness [6]. Automation in organizations has partially preserved this agency by requiring human initiation and input while maintaining supervisory oversight, as well as allowing humans to intervene when necessary [12].

While concerns about sacrificing meaningful human control already exist in traditional automation, agentic AI systems challenge

human agency further by coordinating multiple agents to solve complex organizational tasks with minimal human involvement [5]. When multiple specialized agents coordinate behind unified interfaces to complete work tasks, humans may find themselves reduced to approval roles with limited understanding of how problems were approached or solved [23]. This creates situations where humans maintain formal authority over outcomes but lose meaningful participation in the problem-solving process, potentially leading to a decreased sense of agency even though they retain responsibility for outcomes. Research questions include:

- How do humans maintain meaningful agency when multiple agents coordinate to solve complex tasks with minimal human input?
- What level of human involvement preserves appropriate agency while enabling effective multi-agent problem-solving?
- How does reduced participation in problem-solving processes affect human professional identity and expertise?
- What interface designs support human agency when multi-agent coordination handles complex tasks?

### 2.3 Sustaining Human Skills in Long-Term Human-AI Interaction

Furthermore, AI systems introduce temporal dynamics in interactions at the workplace: workers must continuously adapt to evolving AI capabilities and calibrate their reliance on these systems [10]. Maintaining active involvement in problem-solving and decision-making allows humans to develop expertise alongside AI, fostering collaborative relationships where responsibility and contribution are balanced [30]. However, appropriate reliance is crucial—overreliance or underreliance on AI advice can affect the development and retention of human skills over time [13, 25].

Agentic AI systems create fundamentally different temporal dynamics, and the combination of formal authority with reduced agency threatens to deskill workers over time. When humans maintain decision-making responsibility while multi-agent systems handle complex problem-solving with minimal human involvement, humans may gradually lose the expertise necessary to meaningfully evaluate AI advice or intervene effectively when problems arise [22]. This creates a situation where agentic AI systems lead to a decrease in human competence, making humans increasingly unable to provide the oversight they remain formally responsible for [23]. Over time, efficient interaction arrangements evolve into dependency relationships where humans lack the knowledge to evaluate AI solutions critically or the skills to resume direct involvement in organizational processes. Research questions include:

- How do human skills and expertise change over time when agentic AI handles complex problem-solving with minimal human involvement?
- What temporal patterns characterize the evolution from human oversight to human dependency in agentic AI relationships?
- How can organizations maintain human competence while enabling long-term agentic AI collaboration?
- What intervention strategies prevent deskilling while preserving the benefits of agentic AI autonomy?

## 3 Organizers

The organizing committee possesses comprehensive knowledge spanning user-centered automation and AI research domains. The members have facilitated workshops addressing diverse themes such as user experience, everyday automation, transportation systems, safety-critical contexts as well as human-AI collaboration (e.g., CHI, MobileHCI, AutomotiveUI). Additionally, the team comprises individuals who have held leadership roles, chairing workshops and conferences (e.g., CHI, MobileHCI, HRI, AutomotiveUI, Persuasive), demonstrating their established expertise and standing within the community.

- **Philipp Spitzer** (main contact) is a research lab lead in Information Systems at the Karlsruhe Institute of Technology. He leads the Applied AI in Service Systems lab and conducts research projects on applied machine learning topics. His main focus is on the interplay of knowledge and human-AI interaction, focusing on designing robust and effective collaborative systems.
- **Matthias Baldauf** is a Professor for Business Informatics at the Eastern Switzerland University of Applied Sciences in St.Gallen. He leads user-centered projects in the domain of human-automation interaction in smart manufacturing and office workplaces.
- **Philippe Palanque** is a Professor of Computer Science at the University of Toulouse III. His research focuses on interactive systems design, development, certification, and deployment in various safety critical contexts (e.g., aircraft cockpits, satellite workstations).
- **Virpi Roto** is a Senior University Lecturer and the Academic Leader of the EIT Culture and Creativity at Aalto University. She studies the means to design automation that improves employee experience, for example, in maritime and industrial contexts.
- **Katelyn Morrison** is a PhD student at the Human-Computer Interaction Institute at Carnegie Mellon University who has worked on human-centered agentic AI systems over the last two years with IBM Research. Her research broadly focuses on human-centered AI and human-AI collaboration.
- **Garoa Gomez-Beldarrain** is a PhD student in the Knowledge and Intelligence Design research group at Delft University of Technology. Her research focuses on responsible automation adoption in organizational settings, which she pursues through action research in a major European airport.
- **Monika Westphal** is a Professor of Operations Management at IE University in Madrid. Her research addresses behavioral questions in operations management, with the goal of improving performance and efficiency in service and production processes, while ensuring the well-being of workers and other stakeholders.
- **Joshua Holstein** is a postdoctoral researcher at the Digital Service Innovation chair at the Karlsruhe Institute of Technology. His work focuses on enhancing human-AI collaboration by supporting the development of humans' mental models of AI systems and their underlying data, examining their impact on reliance behavior and team performance.

## 4 Website

All relevant information will be available at <https://matthiasbaldauf.com/automationxp26>, including the workshop’s aims and scope, submission guidelines with recommended topics, the program timeline, details on how to participate, and organizer biographies.

## 5 Pre-Workshop Plans

We will promote the workshop through key HCI mailing lists (including CHI Announcements, Ubicomp Announcements, and others) and relevant online platforms. Additionally, we will send direct invitations to our academic network, encompassing previous workshop participants and organizers of related workshops. We will solicit position papers of up to four pages (references excluded), formatted according to ACM conference guidelines. Submissions should present the authors’ contributions to the workshop themes and may take the form of ongoing research, research proposals, or demonstrations that address relevant research questions within the scope of automation engagement in agentic AI.

The organizing team will evaluate submissions considering their alignment with the workshop focus, novelty, impact, and overall quality. Our target is to accept approximately 25 submissions, which will be made accessible via the workshop website ahead of the event.

## 6 Accessibility

Accepted position paper authors will be required to make their submissions accessible. The organizers will support authors during final submission preparation by providing instructions on document tagging, adding alternative text for images, setting tab orders, and additional best practices. Moreover, we are committed to ensuring the workshop is accessible to all participants and will arrange for real-time captioning, adjusted according to attendee needs.

## 7 Asynchronous Engagement

Asynchronous participation for individuals facing technical or accessibility challenges will be facilitated through various activities both before and after the workshop. One week prior to the event, participants’ position papers will be made available on the workshop website, enabling participants to familiarize themselves with all accepted contributions in advance. This approach allows those with limited English proficiency to prepare ahead of time. Following the workshop, all materials generated during the event will be published on the website for asynchronous access.

## 8 Workshop Format and Structure

The workshop format and organization draw upon lessons learned and proven approaches from prior related workshops [4, 18, 32]. We have planned a two-session in-person workshop with a total duration of 3 hours. The format emphasizes structured discussions within a formal yet collaborative environment to facilitate meaningful engagement for all participants. An overview of the tentative schedule is presented in Table 1.

### 8.1 Session 1

The first session of the workshop comprises an introduction to the workshop theme, presentations by participants, and discussions.

**8.1.1 Introduction.** The workshop will begin with an official welcome to the participants, during which the organizers will present the theme of the workshop and outline its main objectives. This will be followed by brief introductions from the organizers and the participants.

**8.1.2 Paper Madness I.** Prior to the workshop, the organizers will categorize and arrange participant contributions based on their thematic relation. During this first “paper madness” session, around 10 participants will briefly present their recent research. Presentations will be limited to 3-minute and single-slide pitches based on a provided template. This format was intentionally chosen to promote contribution-oriented presentations, facilitate easy reference to participants’ work in subsequent sessions, and provide an interactive medium for contributions.

Additionally, the audience will be able to leave comments, questions, and suggestions during the presentations. We will prepare a shared document for joint note-taking; these notes are supposed to provide inspiration for the following group work and the advancement of the position papers by the authors.

**8.1.3 Discussion - Problem and Solution Spaces.** After the first set of presentations, participants will be placed in smaller groups. The groups will discuss different perspectives on the challenges presented in the workshop regarding agentic automation experience, as well as their manifestations in the context of the participants’ research. Afterwards, they will ideate potential solutions that could be proposed through the application of HCI knowledge and methods. Every group will be moderated by a member of the organizing team, who will guide participant discussions by helping them engage in different brainstorming, reflection, ideation, and definition activities, by combining divergent and convergent facilitation. To facilitate discussion and collaboration, physical materials, such as paper sheets, colored pens, and post-its, will be provided for the groups.

### 8.2 Session 2

The second session will include the second set of paper presentations, followed by further collaborative group activities, a presentation of the outcomes, and a subsequent plenary discussion. The session will conclude with a summary provided by the organizers, where we expect to provide a visual of the resulting research agenda.

**8.2.1 Paper Madness II.** A second set of participants will present their one-slide pitch. We will introduce the participants based on their thematic alignment to make the session easier to follow.

**8.2.2 Discussion - Agentic Automation Experiences research agenda.** Participant groups will identify common themes, research questions, and methodological approaches for designing, assessing, and measuring agentic automation experiences, based on the problem and solution spaces that were previously identified.

This session will also be facilitated by the members of the organizing team. The outcomes of these group activities may include new research ideas and draft proposals for potential (joint) research projects.

**Table 1: Preliminary schedule for an on-site workshop with two 90-minute sessions.**

Duration	Activity (details in text)
<b>Session 1</b>	
15 minutes	<b>Introduction:</b> Welcome, introduction, and activities overview
30 minutes	<b>Paper madness I:</b> First set of participant presentations
45 minutes	<b>Discussion - the problem space and potential solutions:</b> Collaborative discussion and ideation on workshop challenges and potential solutions
<b>Break</b>	
<b>Session 2</b>	
30 minutes	<b>Paper madness II:</b> Second set of participant presentations
30 minutes	<b>Discussion - Agentic Automation Experiences research agenda:</b> Collaborative discussion and ideation of potential research directions for future HCI works
30 minutes	<b>Synthesis:</b> presentation of key results and actions on follow-up activities

*8.2.3 Workshop synthesis.* Towards the end of the workshop, all groups will reconvene in a plenary discussion session to present the outcomes of their discussions. The organizers will then conclude the session by outlining the Agentic Automation Experiences research agenda, a roadmap for future research and initiatives, based on the ideas generated. Efforts will be made to identify relevant research communities, outlining opportunities, timelines, and project formats.

## 9 Post-Workshop Plans

All workshop materials and results will be published on the official workshop website, including photos, key findings, and plans for subsequent activities. To maximize outreach and guarantee archival access, we intend to publish the proceedings through arXiv. Participant consent will be obtained prior to capturing images and sharing workshop materials online. Additionally, we will reach out to appropriate venues to explore opportunities for a special journal issue or magazine contribution featuring the workshop’s submissions and results. Potential targets include journals such as “Personal and Ubiquitous Computing” [17] or “Computer Supported Cooperative Work” [1]. Additionally, based on workshop participation and discussions, we aim to establish an ongoing research community and facilitate collaborations to generate tangible outputs and extend our discussions of this workshop to other venues such as MobileHCI, UbiComp, IUI, and UMAP.

## 10 Call for Participation

Agentic AI has rapidly gained traction and is increasingly integrated across organizations: it shifts automation experiences from specific tools to proactive systems coordinated by multiple specialized agents, which are able to solve complex problems with minimal human involvement. This workshop aims to examine the critical challenges posed by agentic AI in organizational settings, particularly regarding enabling multi-agent transparency and attribution, balancing human agency in agentic AI, and sustaining human skills in long-term human-AI interaction.

We invite position papers based on recent or forthcoming research, case studies, or design work. We encourage participants to share

insights from real-world implementations and challenges of multi-agent coordination, agency, and evolving human-AI relationships. Relevant topics of interest include, but are not limited to:

- Attribution mechanisms for multi-agent coordination
- Balancing human agency with autonomous multi-agent problem-solving
- Understanding when reduced agency improves versus harms human well-being and outcomes
- Maintaining human expertise and preventing deskilling in long-term agentic AI collaboration
- Real-world organizational cases of agentic AI adoption
- Socio-technical considerations for integrating agentic AI into workflows

Submissions should follow the single-column ACM conference proceedings format and may not exceed four pages (excluding references). Non-anonymized papers must be submitted via <https://easychair.org/conferences/?conf=automationxp26>. The organizing committee will evaluate submissions based on their relevance, originality, significance, and quality. At least one author of each accepted paper is required to attend the workshop. Important dates:

- Position paper deadline: February 13th, 2026
- Acceptance notification: February 20th, 2026
- Camera-ready versions due: March 27th, 2026
- Workshop date: April 13th, 2026

Website: <https://matthiasbaldauf.com/automationxp26>

## 11 Expected Attendance

We expect around 30-40 registered on-site attendees.

## References

- [1] 2025. Automation Experiences – Communication, Coordination, and Collaboration within Human-AI Teams. Special issue announcement, Computer Supported Cooperative Work (CSCW). Open for submissions; submission deadline: 31 October 2025.
- [2] Deepak Bhaskar Acharya, Karthigeyan Kuppan, and B Divya. 2025. Agentic ai: Autonomous intelligence for complex goals—a comprehensive survey. *IEE Access* (2025).
- [3] L Bainbridge. 1983. Ironies of Automation. *Automatica* 19 (1983), 775–779. Issue 6.
- [4] Matthias Baldauf, Peter Fröhlich, Virpi Roto, Philippe Palanque, Siân Lindley, Jon Rogers, Wendy Ju, and Manfred Tscheligi. 2022. Engaging with Automation: Understanding and Designing for Operation, Appropriation, and Behaviour Change. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI EA '22). Association for Computing Machinery, New York, NY, USA, Article 73, 6 pages. doi:10.1145/3491101.3503735
- [5] Ajay Bandi, Bhavani Kongari, Roshini Naguru, Sahitya Pasnoor, and Sri Vidya Vilipala. 2025. The Rise of Agentic AI: A Review of Definitions, Frameworks, Architectures, Applications, Evaluation Metrics, and Challenges. *Future Internet* 17, 9 (2025), 404.
- [6] Albert Bandura. 2006. Toward a psychology of human agency. *Perspectives on psychological science* 1, 2 (2006), 164–180.
- [7] Anjanava Biswas and Wrick Talukdar. 2025. *Building Agentic AI Systems: Create intelligent, autonomous AI agents that can reason, plan, and adapt*. Packt Publishing Ltd.
- [8] Christian Schroeder de Witt. 2025. Open challenges in multi-agent security: Towards secure systems of interacting ai agents. *arXiv preprint arXiv:2505.02077* (2025).
- [9] Alessandro Delfanti and Bronwyn Frey. 2021. Humanly Extended Automation or the Future of Work Seen through Amazon Patents. *Science, Technology, & Human Values* 46, 3 (2021), 655–682. arXiv:https://doi.org/10.1177/0162243920943665 doi:10.1177/0162243920943665
- [10] Dominik Dellermann, Philipp Ebel, Matthias Söllner, and Jan Marco Leimeister. 2019. Hybrid intelligence. *Business & Information Systems Engineering* 61, 5 (2019), 637–643.
- [11] Zhuoyun Du, Chen Qian, Wei Liu, Zihao Xie, Yifei Wang, Rennai Qiu, Yufan Dang, Weize Chen, Cheng Yang, Ye Tian, et al. 2025. Multi-Agent Collaboration via Cross-Team Orchestration. In *Findings of the Association for Computational Linguistics: ACL 2025*. 10386–10406.
- [12] Mica R Endsley. 2017. From here to autonomy: lessons learned from human-automation research. *Human factors* 59, 1 (2017), 5–27.
- [13] Avigail Ferdman. 2025. AI, Desksilling, and the Prospects for Public Reason. *Minds and Machines* 35, 3 (2025), 1–27.
- [14] Adam Fourney, Gagan Bansal, Hussein Mozannar, Cheng Tan, Eduardo Salinas, Friederike Niedtner, Grace Proebsting, Griffin Bassman, Jack Gerrits, Jacob Alber, et al. 2024. Magentic-one: A generalist multi-agent system for solving complex tasks. *arXiv preprint arXiv:2411.04468* (2024).
- [15] Sarah E. Fox and Samantha Shorey. 2024. *Artfully Integrating AI: Proceeding Responsibly With Worker-Centered Best Practices*. Springer International Publishing, Cham, 67–84. doi:10.1007/978-3-031-66528-8\_4
- [16] Sarah E. Fox, Samantha Shorey, Esther Y. Kang, Dominique Montiel Valle, and Estefania Rodriguez. 2023. Patchwork: The Hidden, Human Labor of AI Integration within Essential Work. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW1, Article 81 (apr 2023), 20 pages. doi:10.1145/3579514
- [17] Peter Fröhlich, Matthias Baldauf, Thomas Meneweger, Manfred Tscheligi, Boris de Ruyter, and Fabio Paternó. 2020. Everyday automation experience: a research agenda. *Personal and Ubiquitous Computing* 24, 6 (Oct. 2020), 725–734. doi:10.1007/s00779-020-01450-y
- [18] Peter Fröhlich, Matthias Baldauf, Philippe Palanque, Virpi Roto, Fabio Paternó, Wendy Ju, and Manfred Tscheligi. 2023. Intervening, Teaming, Delegating: Creating Engaging Automation Experiences. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI EA '23). Association for Computing Machinery, New York, NY, USA, Article 357, 6 pages. doi:10.1145/3544549.3573799
- [19] Garoa Gomez-Beldarrain, Himanshu Verma, Euiyoung Kim, and Alessandro Bozzon. 2025. Why does Automation Adoption in Organizations Remain a Fallacy?: Scrutinizing Practitioners' Imaginaries in an International Airport. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 213, 19 pages. doi:10.1145/3706598.3713978
- [20] Christopher Gyldenkerne, Jens Ulrik Hansen, Morten Hertzum, and Troels Mønsted. 2024. Innovation tactics for implementing an ML application in healthcare: A long and winding road. *International Journal of Human-Computer Studies* 181 (2024), 103162. doi:10.1016/j.ijhcs.2023.103162
- [21] Mathew Iantorno, Olivia Doggett, Priyank Chandra, Julie Yujie Chen, Rosemary Steup, Noopur Raval, Vera Khovanskaya, Laura Lam, Anubha Singh, Sarah Rotz, and Matt Ratto. 2022. Outsourcing Artificial Intelligence: Responding to the Reassertion of the Human Element into Automation. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI EA '22). Association for Computing Machinery, New York, NY, USA, Article 103, 5 pages. doi:10.1145/3491101.3503720
- [22] Burak Koçak and İsmail Meşe. 2025. AI agents in radiology: toward autonomous and adaptive intelligence. *Diagnostic and Interventional Radiology* (2025).
- [23] Paul M. Leonardi. 2025. Homo agenticus in the age of agentic AI: Agency loops, power displacement, and the circulation of responsibility. *Information and Organization* 35, 3 (2025), 100582. doi:10.1016/j.infoandorg.2025.100582
- [24] Yaxi Lu, Shenzhi Yang, Cheng Qian, Guirong Chen, Qinyu Luo, Yesai Wu, Huadong Wang, Xin Cong, Zhong Zhang, Yankai Lin, et al. [n. d.]. Proactive Agent: Shifting LLM Agents from Reactive Responses to Active Assistance. In *The Thirteenth International Conference on Learning Representations*.
- [25] Katelyn Morrison, Philipp Spitzer, Violet Turri, Michelle Feng, Niklas Kühl, and Adam Perer. 2024. The Impact of Imperfect XAI on Human-AI Decision-Making. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW1 (2024), 1–39.
- [26] Suchismita Naik, Amanda Snellinger, Austin L Toombs, Scott Saponas, and Amanda K Hall. 2025. Exploring Early Adopters' Use of AI Driven Multi-Agent Systems to Inform Human-Agent Interaction Design: Insights from Industry Practice. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 1–8.
- [27] Venkata Siva Prakash Nimmagadda. 2019. Explainable AI in Regulatory Compliance for Pharmaceutical Manufacturing. *European Journal of Quantum Computing and Intelligent Agents* 3 (2019), 341–382.
- [28] Inoluwa Deborah Raji, Andrew Smart, Rebecca N White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 33–44.
- [29] Simone Rozzi and Paola Amaldi. 2012. Organizational and inter-organizational precursors to problematic automation in safety critical domains. In *ATACCS '12: Proceedings of the 2nd International Conference on Application and Theory of Automation in Command and Control Systems*. Association for Computing Machinery, New York, NY, USA, 98–106.
- [30] Max Schemmer, Andrea Bartos, Philipp Spitzer, Patrick Hemmer, Niklas Kühl, Jonas Liebschner, and Gerhard Satzger. 2023. Towards effective human-AI decision-making: The role of human learning in appropriate reliance on AI advice. *arXiv preprint arXiv:2310.02108* (2023).
- [31] Max Schemmer, Niklas Kuehl, Carina Benz, Andrea Bartos, and Gerhard Satzger. 2023. Appropriate Reliance on AI Advice: Conceptualization and the Effect of Explanations. In *Proceedings of the 28th International Conference on Intelligent User Interfaces* (Sydney, NSW, Australia) (IUI '23). Association for Computing Machinery, New York, NY, USA, 410–422. doi:10.1145/3581641.3584066
- [32] Philipp Spitzer, Matthias Baldauf, Philippe Palanque, Virpi Roto, Katelyn Morrison, Dominique Zipperling, and Joshua Holstein. 2025. Hybrid Automation Experiences – Communication, Coordination, and Collaboration within Human-AI Teams. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (CHI EA '25). Association for Computing Machinery, New York, NY, USA, Article 789, 6 pages. doi:10.1145/3706599.3706738
- [33] Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, et al. 2024. Autogen: Enabling next-gen LLM applications via multi-agent conversations. In *First Conference on Language Modeling*.
- [34] Hubert D. Zając, Dana Li, Xiang Dai, Jonathan F. Carlsen, Finn Kensing, and Tariq O. Andersen. 2023. Clinician-Facing AI in the Wild: Taking Stock of the Sociotechnical Challenges and Opportunities for HCL. *ACM Trans. Comput.-Hum. Interact.* 30, 2, Article 33 (March 2023), 39 pages. doi:10.1145/3582430