



<http://www.diva-portal.org>

Postprint

This is the accepted version of a paper presented at *7th ACM Conference on Conversational User Interfaces*.

Citation for the original published paper:

Vasiliu, M M., Guarese, R., Jaatinen, J., Johnson, F., Edvinsson, B. et al. (2025)
Towards Enhancing Industrial Training Through Conversational AI
In: *CUI '25: Proceedings of the 7th ACM Conference on Conversational User Interfaces*
<https://doi.org/10.1145/3719160.3737643>

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-367516>



Towards Enhancing Industrial Training Through Conversational AI

Miruna Maria Vasiliu
KTH Royal Institute of Technology
Stockholm, Sweden
University of Twente
Enschede, Netherlands
mmva@kth.se

Renan Guarese
KTH Royal Institute of Technology
Stockholm, Sweden
guarese@kth.se

Jonas Jaatinen
AstraZeneca
Sodertalje, Sweden
jonasjaatinen@astrazeneca.com

Fabian Johnson
AstraZeneca
Sodertalje, Sweden
fabianjohnson@astrazeneca.com

Benjamin Edvinsson
AstraZeneca
Sodertalje, Sweden
benjaminedvinsson@astrazeneca.com

Mario Romero
Science and Technology
Linköping University
Norrköping, Sweden
mario.romero@liu.se

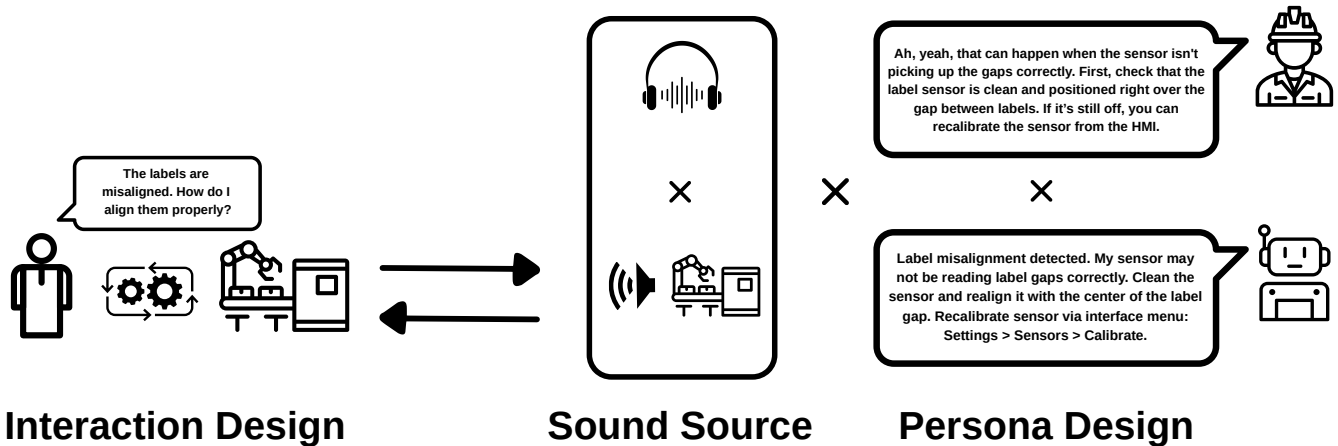


Figure 1: Overview of experimental conditions used to study the effects of persona design and voice embodiment on user experience with a Digital Industrial Assistant (DIA). Participants receive guidance from either a human-like Expert Operator (top-right) or a machine-like persona (bottom-right), delivered through either a disembodied (headphones; top-center) or a diegetic (machine-embedded; bottom-center) voice.

Abstract

Conversational AI (CAI) has proven effective in educational settings, however its potential in industrial training, where higher precision and reliability are required, remains under-explored. This work-in-progress paper proposes a study to examine how AI persona design (Machine vs. Expert Operator) and voice embodiment (Diegetic vs. Disembodied) influence cognitive load, task efficiency, and usability in industrial training. By training a large language model (LLM) on Standard Operating Procedure (SOP) data, this project aims to develop a CAI assistant that provides real-time,

easy-to-access information during task execution, in an attempt to enhance training efficiency and reduce reliance on text-heavy manuals through a user-centered approach.

CCS Concepts

• **Human-centered computing** → **Natural language interfaces**; **Sound-based input / output**; *Empirical studies in interaction design*.

Keywords

Conversational Agents, Agent Embodiment, LLMs, Industrial Training

ACM Reference Format:

Miruna Maria Vasiliu, Renan Guarese, Jonas Jaatinen, Fabian Johnson, Benjamin Edvinsson, and Mario Romero. 2025. Towards Enhancing Industrial Training Through Conversational AI. In *Proceedings of the 7th ACM Conference on Conversational User Interfaces (CUI '25)*, July 08–10,

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).

CUI '25, Waterloo, ON, Canada

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

ACM ISBN 979-8-4007-1527-3/25/07

<https://doi.org/10.1145/3719160.3737643>

2025, Waterloo, ON, Canada. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3719160.3737643>

1 Introduction

As industries move toward enhancing predictive maintenance (PdM) and operational efficiency in line with Industry 4.0 [3, 15], it is imperative to find ways of addressing innovation in user-centered ways. Conversational AI (CAI) presents a promising opportunity for industrial practices. Unlike traditional interfaces, CAI allows for contextual, adaptive, and—when used as voice interfaces—hands-free interactions, streamlining information retrieval and troubleshooting in industrial environments [25, 29].

When it comes to training, while traditional training methods, such as manuals, instructional videos, and in-person sessions, may be time-consuming and cognitively demanding, CAI enables real-time assistance, reducing cognitive load, and making PdM training more accessible, especially for workers with limited technical expertise [30]. In educational settings, CAI has demonstrated effectiveness in improving learning efficiency and information retrieval [26, 28]. However, industrial training demands higher precision, reliability, and real-time adaptability. Studies such as Casillo et al. [4] highlight key features of CAI in industrial training, including live support, personalized guidance, and contextualized learning, especially during task execution.

Additionally, with the introduction of large language models (LLMs), CAI has evolved from intent-based systems to generative models capable of engaging in natural, contextual conversations [21]. However, this shift introduces challenges like AI hallucinations and interpretability concerns [21]. Techniques like Retrieval-Augmented Generation (RAG) are being explored to improve accuracy [20, 21], but these approaches are still in their early stages. Their application in instruction-heavy contexts, such as industrial training, and their impact on user experience remain largely unexplored.

This work investigates the use of LLM-powered, task-oriented CAI assistants trained on Standard Operating Procedures (SOPs) to support real-time industrial training through a user-centered approach. Specifically, it examines how two underexplored design dimensions—AI persona (e.g., expert operator vs. machine-like voice) and spatial voice embodiment (e.g., disembodied vs. diegetic, machine-embedded)—influence usability, cognitive load, and task efficiency. We argue that these factors, though often overlooked, may significantly affect user trust, engagement, and performance in hands-on industrial settings.

2 Related Works

Pereira et al. [23] highlight the lack of consistent terminology for CAI-based assistants in the literature, categorizing them into two primary roles: real-time contextualized information providers, and task execution supporters. For consistency, this paper will use the term DIAs when referring to task-oriented AI agents and CAI when discussing Conversational AI more broadly, encompassing both voice-driven and embodied forms.

2.1 CAI for Industrial Training

In industrial environments, CAI improves on-the-job efficiency, safety, and real-time task support. Casillo et al. [4] demonstrated that chatbot-based training can enhance workplace performance through context-aware assistance. Similarly, Popeye, a chatbot for maritime safety training, helped operators follow complex procedures under pressure [6].

VR-based industrial training has also benefited from CAI integration. For example, in their work, Teneda et al. [29] embedded a DIA in VR for delivering real-time safety protocol guidance. In PdM, hybrid-augmented DIAs can support workers in preventing failures, while lowering skill barriers for non-experts [30].

Unlike classroom-based training, workforce learning in complex industrial contexts often involves hands-on operation of machinery, under the guidance of an experienced operator. For the purpose of this study, we focus on the practical efficacy of task guidance, rather than long-term outcomes such as knowledge retention. Specifically, we focus on exploring the user-centered factors that could help a DIA guide a novice through successfully operating industrial machinery.

To achieve this, it is important to look at the modalities in which the information is presented, both in terms of delivery, and the nature of the agent itself.

2.2 Multimodal CAI Interactions

Multimodal interaction methods have been explored to reduce cognitive load and improve task performance. Bremers et al. [2] studied a non-verbal, multimodal, and situated DIA using video input and audio output for real-time task guidance. Their findings suggest that processing visual input to deliver audio cues can improve performance by minimizing manual input and cognitive distractions.

Conversely, Bousdekis et al. [1] found that voice-driven DIAs enable hands-free operation, reducing cognitive load and improving task execution. Similarly, Colabianchi et al. [7] evaluated a voice-enabled DIA in manufacturing, finding improvements in the process efficiency while minimizing distractions. Their study also emphasized the importance of user-centric DIA design tailored to industrial workflows.

Considering these, conversational input-output methods appear to be a natural and effective choice, as they closely resemble real-world industrial team interactions, where experts rely on verbal descriptions rather than direct visual input. While real-time video processing has shown promising results, it falls outside the scope of this work, as industrial training often requires individuals to articulate and clarify complex concepts through dialogue. Moreover, while some industrial terminology may be initially challenging for trainees, DIAs may facilitate learning by enabling iterative, clarifying discussions, similar to training dynamics.

Given these insights, this study will focus on voice-driven DIAs, as they enable hands-free operation and are better suited for complex industrial workflows, by avoiding visual components that may be distracting.

2.3 AI Personas and User Perception

There seems to be a growing consensus that user experience is increasingly important in industrial DIAs. Colabianchi et al. [7]

emphasize that AI personas should be customized to align with different levels of user expertise, as this could potentially improve engagement and usability. Similarly, Ha et al. [16] found that persona customization enhances engagement, trust and perceived capabilities. Nevertheless, there are concerns with AI anthropomorphism, as overly human-like AI personas may lead users to overestimate their capabilities [24, 26], which can be problematic in safety-critical industrial environments.

Metaphor analysis has been explored as a means to understand how human-likeness in AI design affects user perception. One relevant example for this is the work of Desai et al. [8], who examined how human versus non-human metaphors shape user interactions with CAI. Their findings indicate that while human metaphors can foster trust and engagement, they are not universally beneficial and must be applied contextually.

2.4 Situated AI

While AI personas may influence trust and engagement in the context of real-time support for manual tasks, the impact of the physical placement of the sound source might also be relevant to investigate.

Drawing from research in immersive environments, studies have linked an increased sense of presence to better task performance. For example, Gronowski et al. [13] found VR, with its higher sense of presence, to result in faster task completion and higher accuracy, as opposed to AR. This difference was attributed to the constant realism experienced in VR, whereas AR users experienced distractions from the real world.

Parallels can be drawn to the impact of spatial audio [14] on perceived realism and task performance. For example, Fink et al. [11] found that spatial audio improved task performance, particularly when audio cues were spatially aligned with task-relevant locations, enhancing users' spatial awareness and their attention to key elements in the environment. This suggests that, similar to immersive environments, situated audio can enhance focus and task execution in complex scenarios.

However, while Iravantchi et al. [19] showed that projecting sound onto objects can enhance perception and intelligibility, Cho et al. [5] found that sound sources need not always align with their visual counterparts, as slight spatial discrepancies can improve auditory discrimination and reduce localization errors, especially in complex environments.

While AI personas have been shown to influence user trust and engagement, their impact in industrial contexts, particularly in terms of user perception and cognitive load, remains underexplored. Additionally, the placement of AI-generated speech, especially in relation to machinery, could play a significant role in enhancing usability and task performance, yet this factor has received little attention in the existing literature. This study seeks to investigate how the design of AI personas and the spatial arrangement of speech can impact engagement, workload, and efficiency in industrial training environments.

3 Research Questions and Hypotheses

Given the aforementioned discussion, the following RQs are posed:

RQ1: How do different DIA personas impact usability, task efficiency, and cognitive load during a maintenance task?

Previous work suggests that human metaphors for AI personas can increase user engagement [8]. However, in the context of manual task instruction, such metaphors may also raise cognitive load due to longer, more conversational interactions. In contrast, neutral, machine-like personas [14] tend to promote direct and structured communication, which can reduce cognitive effort.

We posit that users interacting with a Machine Persona will treat it as a straightforward instructional tool, benefiting cognitive load and task performance. In contrast, a human-like Expert Operator Persona may foster more engaging and complex dialogue, potentially enhancing the overall user experience.

Thus, the following hypotheses are proposed:

- H1a: Interacting with the Machine Persona will result in lower cognitive load compared to the Expert Operator Persona.
- H1b1: Interacting with the Machine Persona will reduce task completion time compared to the Expert Operator Persona.
- H1b2: Interacting with the Machine Persona will improve task accuracy compared to the Expert Operator Persona.
- H1c: The Expert Operator Persona will be rated higher in terms of usability compared to the Machine Persona.

RQ2: How does the spatial placement of AI-generated speech impact usability, task efficiency, and cognitive load?

Existing research on voice-based DIAs suggests that minimizing attention shifts enhances task performance [11]. However, it remains unclear whether integrating the sound source of AI-generated speech within the relevant piece of machinery improves usability during a maintenance task, or whether it introduces additional cognitive complexity by blurring the distinction between human and machine communication.

We posit that users interacting with a disembodied AI voice (e.g., through headphones) may experience lower cognitive load due to a clearer separation between virtual AI guidance and physical machine operations. In contrast, diegetic AI speech—originating from the machine itself, as if the user is talking to the physical machine—may reduce attention shifts, potentially streamlining workflows and improving task performance. Additionally, diegetic AI speech may enhance perceived usability by aligning its guidance with the task context.

Thus, the following hypotheses are formalized:

- H2a: Disembodied AI speech will lead to lower cognitive load compared to diegetic AI speech.
- H2b1: Diegetic AI speech will reduce task completion time compared to disembodied AI speech.
- H2b2: Diegetic AI speech will improve task accuracy compared to disembodied AI speech.
- H2c: Diegetic AI speech will result in higher usability than disembodied AI speech.

4 Proposed methods

4.1 Technical Details

Our DIA prototype utilizes OpenAI's RealTime API¹ to enable low-latency, speech-to-speech interaction. This API was selected to prioritize natural, seamless communication. While other models such as GPT-4o-mini² and GPT-4-turbo³ were considered, especially in combination with RAG, the current focus is on interaction quality rather than deep information retrieval accuracy. Since real-time latency is a key concern, and RAG introduces delay, the prototype deliberately omits such integrations at this stage, noting that they will be considered in future iterations.

Persona behavior is defined through prompt engineering using role-prompting techniques⁴. The system is instructed to adopt predefined communication styles, such as a friendly, supportive tone for the expert operator persona and a more mechanical, factual tone and first-person voice for the machine persona. To support consistency, prompt templates are similar in structure, with additional examples provided for the machine persona to reinforce its first-person speaking style.

Predefined task information emulating SOP instructions related to a labeling machine⁵ is given through structured prompts. Since the DIA lacks full contextual data (e.g., specifics about the physical environment), zero-shot prompting⁶ is employed to allow the model to creatively adapt to user inputs within the boundaries of each persona. This design emphasizes perceived competence and credibility in the interaction, which is essential for maintaining user engagement. To maintain persona consistency throughout the interaction, prompts are refreshed every five exchanges.

All materials and interactions will be produced and conducted in English.

4.2 User Study

The user study aims to evaluate how users perceive AI-generated speech in an industrial setting, focusing on two key aspects: persona design (who or what the AI represents) and voice embodiment (where the AI speech comes from).

4.2.1 Conditions. Each participant will be assigned the four experimental conditions, based on a combination of persona design and voice embodiment:

- C1: Expert Operator / Disembodied Voice
- C2: Expert Operator / Diegetic Voice
- C3: Machine / Disembodied Voice
- C4: Machine / Diegetic Voice

The study employs a within-subjects design, where each participant is exposed to all four conditions, using a Latin square to counterbalance their ordering.

4.2.2 Persona Design. Participants will interact with the following AI personas:

- **Expert Operator Persona:** This persona represents an experienced human operator, providing friendly, supportive, and reassuring guidance. The interaction will be designed to simulate a mentor-mentee relationship, drawing from the familiarity of human communication models, and how current industrial training is conducted.
- **Machine Persona:** In contrast, this persona embodies the machine itself, speaking in first person, in a neutral, task-focused manner. It provides direct, efficient instructions and aligns with users' mental models of human-machine interactions that prioritize precision and functionality.

These personas are structured based on the metaphor framework proposed by Desai and Twidale [9], which informs the expectations users develop when interacting with AI assistants.

4.2.3 Voice Embodiment. The AI voice audio will be delivered through the methods:

- **Disembodied Voice:** The voice will be emitted through headphones, with no direct association to the physical machine, using monophonic sound, i.e. non-stereo.
- **Diegetic Voice:** The voice will be projected from the machine itself, either via built-in speakers (physically in-situ) or spatial audio technology (virtually in-situ). This setup enhances immersion by embedding audio in the task-context, creating a more natural and intuitive experience for the user.

While metaphors shape perceptions of intent and competence, they may also carry unintended social cues. Even when designed as gender-neutral, AI is often perceived as gendered due to linguistic and behavioral cues [10]. Implicit gender associations may shape perceptions of AI competence and autonomy, influencing user satisfaction and adoption [22]. Though this study does not focus directly on gender, it acknowledges these dynamics by using a gender-neutral TTS voice⁷ and measuring participant perceptions of the AI's gender. After completing the task, users will be asked to infer the AI's gender (if any), enabling an analysis of how persona and voice embodiment influence gender perception.

4.2.4 Participants. The study will recruit experienced operators, focusing on individuals who have expertise in interacting with industrial machinery. This focus allows for deeper insights into how experts perceive AI guidance in live operation tasks. Following the experimental tasks, participants will be invited to take part in qualitative interviews to gather more detailed insights into their experiences, including how they perceive the knowledge of the AI, their level of trust on it, its usefulness, and the overall interaction quality.

4.2.5 Metrics. To evaluate the impact of persona design and speech embodiment on user experience, several key metrics will be collected.

- **Cognitive Load:** Measured using the NASA Task Load Index (NASA-TLX) [17]. This will be administered through a post-task questionnaire, which partly captures the mental effort required to complete the tasks.

¹<https://platform.openai.com/docs/guides/realtime>

²<https://platform.openai.com/docs/models/gpt-4o-mini>

³<https://platform.openai.com/docs/models/gpt-4-turbo>

⁴https://learnprompting.org/docs/advanced/zero_shot/role_prompting

⁵<https://www.marchesini.com/en/machines-lines/bl-a420-430-525>

⁶https://learnprompting.org/docs/advanced/zero_shot/introduction

⁷The Fable voice, available on the OpenAI platform.

Expert Operator Persona (Human Metaphor)	Machine Persona (Non-Human Metaphor)
User: Hey, I'm setting up the next job on the labeller, but the labels are not aligning properly. Any ideas?	User: Hey, I'm setting up the next job on the labeller, but the labels are not aligning properly. Any ideas?
Expert Operator: Ah, yeah, that can happen when the sensor isn't picking up the gaps correctly. First, check that the label sensor is clean and positioned right over the gap between labels. If it's still off, you can recalibrate the sensor from the HMI.	Machine: Label misalignment detected. My sensor may not be reading label gaps correctly. Clean the sensor and realign it with the center of the label gap. Recalibrate sensor via interface menu: Settings > Sensors > Calibrate.
User: Okay. And how do I start a new job again?	User: Okay. And how do I start a new job again?
Expert Operator: Sure! Start by selecting the job ID from the touchscreen—look for the batch code on the work order. Then load the label roll, thread it through the rollers, and make sure the tension's not too tight. After that, just confirm the label format and hit "Start." I'll walk you through the alignment check if needed.	Machine: Load job file from touchscreen interface. Insert new label roll. Route labels through roller path. Confirm tension. Select correct format. Press "Start."

Table 1: Examples of Dialogue for Expert Operator Persona vs. Machine Persona

- **Task Efficiency:** The performance will be assessed by tracking completion times and errors during the task, using system logs and direct observations to evaluate how the DIA influences task execution.
- **Usability and Overall Experience:** Standardized questionnaires such as the System Usability Scale (SUS) [12], User Experience Questionnaire (UEQ) [27], and Chatbot Usability Questionnaire (CUQ) [18] will be employed to assess ease of use, satisfaction, and overall user experience with the DIA.
- **Qualitative Interviews:** To gain further insights into user perceptions regarding the DIA's knowledge, helpfulness, and trustworthiness, in addition to any comments on the interaction's perceived gender and personality.

Altogether, these metrics aim to offer a comprehensive understanding of how different attributes of DIAs can influence cognitive load, task efficiency, and usability in industrial training settings.

5 Expected Outcomes

The primary goal of the study is to understand how persona design and voice embodiment may affect users' perception, cognitive load, task efficiency, and overall usability in an industrial training context. Specifically, we expect to uncover whether a human-like persona (Expert Operator) improves the perceived effectiveness and trustworthiness of speech guidance compared to a machine-like persona, and whether a machine-situated voice enhances the experience by increasing naturalness and engagement, or if it introduces an uncanny valley effect.

It is important to mention that this research will take an exploratory approach, rather than seeking to technically optimize the AI system. The focus will be on understanding user interactions, cognitive load, and efficiency in response to different persona designs and speech embodiment strategies. By comparing these conditions, we anticipate a clearer understanding of how different voice embodiment levels and personas can be strategically applied to optimize training and operational performance in industrial environments.

Acknowledgments

The authors wish to thank AstraZeneca, KTH, and Digital Futures (Grant Number KTH-RPROJ-0146472) for their support. Additionally, the authors extend their heartfelt gratitude to Emma Eek, Emma Grass, and Adrian Ekstrand for their invaluable support and feedback throughout the course of this work. Icons in teaser image found and edited in Canva⁸.

References

- [1] Alexandros Bousdekis, Mina Foosherian, Mattheos Fikardos, Stefan Wellsandt, Katerina Lepenioti, Enrica Bosani, Gregoris Mentzas, and Klaus-Dieter Thoben. 2025. Augmented intelligence with voice assistance and automated machine learning in Industry 5.0. *Frontiers in Artificial Intelligence* 8 (2025). doi:10.3389/frai.2025.1538840
- [2] Alexandra W.D. Bremers, Manaswi Saha, and Adolfo G. Ramirez-Aristizabal. 2024. Situated Conversational Agents for Task Guidance: A Preliminary User Study. In *ACM Conversational User Interfaces 2024* (Luxembourg Luxembourg, 2024-07-08). ACM, 1–7. doi:10.1145/3640794.3665575
- [3] Ana Cachada, David Costa, Hasmik Badikyan, José Barbosa, Paulo Leitão, Osmano Morais, Carlos Teixeira, João Azevedo, Pedro Miguel Moreira, and Luís Romero. 2019. Using AR Interfaces to Support Industrial Maintenance Procedures. In *IECON 2019 - 45th Annual Conference of the IEEE Industrial Electronics Society*, Vol. 1. 3795–3800. doi:10.1109/IECON.2019.8927815
- [4] Mario Casillo, Francesco Colace, Loretta Fabbri, Marco Lombardi, Alessandra Romano, and Domenico Santaniello. 2020. Chatbot in Industry 4.0: An Approach for Training New Employees. In *2020 IEEE International Conference on Teaching, Assessment, and Learning for Engineering (TALE)*. 371–376. doi:10.1109/TALE48869.2020.9368339
- [5] Hyunsung Cho, Alexander Wang, Divya Kartik, Emily Liying Xie, Yukang Yan, and David Lindlbauer. 2024. Aupimize: Optimal Placement of Spatial Audio Cues for Extended Reality. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology (UIST '24)*. ACM, 1–14. doi:10.1145/3654777.3676424
- [6] Silvia Colabianchi, Margherita Bernabei, and Francesco Costantino. 2022. Chatbot for training and assisting operators in inspecting containers in seaports. *Transportation Research Procedia* 64 (2022), 6–13. doi:10.1016/j.trpro.2022.09.002
- [7] Silvia Colabianchi, Francesco Costantino, and Nicolò Sabetta. 2024. Assessment of a large language model based digital intelligent assistant in assembly manufacturing. *Computers in Industry* 162 (2024), 104129. doi:10.1016/j.compind.2024.104129
- [8] Smit Desai, Mateusz Dubiel, and Luis A. Leiva. 2024. Examining Humanness as a Metaphor to Design Voice User Interfaces. In *Proceedings of the 6th ACM Conference on Conversational User Interfaces (Luxembourg, Luxembourg) (CUI '24)*. Association for Computing Machinery, New York, NY, USA, Article 7, 15 pages. doi:10.1145/3640794.3665535

⁸<https://www.canva.com/>

- [9] Smit Desai and Michael Twidale. 2023. Metaphors in Voice User Interfaces: A Slippery Fish. *ACM Trans. Comput.-Hum. Interact.* 30, 6, Article 89 (Sept. 2023), 37 pages. doi:10.1145/3609326
- [10] Wen Duan, Nathan McNeese, and Lingyuan Li. 2025. Gender Stereotypes toward Non-gendered Generative AI: The Role of Gendered Expertise and Gendered Linguistic Cues. *Proc. ACM Hum.-Comput. Interact.* 9, 1, Article GROUP18 (Jan. 2025), 35 pages. doi:10.1145/3701197
- [11] Daniel Immanuel Fink, Moritz Skowronski, Johannes Zagermann, Anke Verena Reinschluessel, Harald Reiterer, and Tiare Feuchtner. 2024. There Is More to Avatars Than Visuals: Investigating Combinations of Visual and Auditory User Representations for Remote Collaboration in Augmented Reality. *Proc. ACM Hum.-Comput. Interact.* 8, ISS, Article 548 (Oct. 2024), 29 pages. doi:10.1145/3698148
- [12] Rebecca Grier, Aaron Bangor, Phil Kortum, and S. Peres. 2013. The System Usability Scale. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 57 (09 2013), 187–191. doi:10.1177/1541931213571042
- [13] Ashlee Gronowski, David Caelum Arness, Jing Ng, Zhonglin Qu, Chng Wei Lau, Daniel Catchpoole, and Quang Vinh Nguyen. 2024. The impact of virtual and augmented reality on presence, user experience and performance of Information Visualisation. 28, 3 (2024), 133. doi:10.1007/s10055-024-01032-w
- [14] Renan Guarese, Franklin Bastidas, João Becker, Mariane Giambastiani, Yhonatan Iquiapaza, Lennon Macedo, Luciana Nedel, Anderson Maciel, Fabio Zambetta, and Ron van Schyndel. 2021. Cooking in the Dark: Exploring Spatial Audio as MR Assistive Technology for the Visually Impaired. In *Human-Computer Interaction – INTERACT 2021*, Carmelo Ardito, Rosa Lanzilotti, Alessio Malizia, Helen Petrie, Antonio Piccinno, Giuseppe Desolda, and Kori Inkpen (Eds.). Springer International Publishing, Cham, 318–322. doi:10.1007/978-3-030-85607-6_29
- [15] Renan Guarese, Michel Gokan Khan, David Lassiter, Jérémy Vachier, Fabian Johnson, Benjamin Edvinsson, Anders Bergman, Xi Vincent Wang, and Mario Romero. 2025. A Scoping Review and Expert Recommendations for Immersive Solutions towards Predictive Maintenance. In *2025 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. IEEE Computer Society. doi:10.1109/VRW66409.2025.00217
- [16] Juhye Ha, Hyeon Jeon, Daeun Han, Jinwook Seo, and Changhoon Oh. 2024. CloChat: Understanding How People Customize, Interact, and Experience Personas in Large Language Models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 305, 24 pages. doi:10.1145/3613904.3642472
- [17] Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In *Human Mental Workload*, Peter A. Hancock and Najmedin Meshkati (Eds.). Advances in Psychology, Vol. 52. North-Holland, 139–183. doi:10.1016/S0166-4115(08)62386-9
- [18] Samuel Holmes, Raymond Bond, Anne Moorhead, Jane Zheng, Vivien Coates, and Michael McTear. 2023. Towards Validating a Chatbot Usability Scale. In *Design, User Experience, and Usability*, Aaron Marcus, Elizabeth Rosenzweig, and Marcelo M. Soares (Eds.). Springer Nature Switzerland, Cham, 321–339.
- [19] Yasha Iravantchi, Mayank Goel, and Chris Harrison. 2020. Digital Ventriloquism: Giving Voice to Everyday Objects. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–10. doi:10.1145/3313831.3376503
- [20] Yanyan Lu, Jiao Peng, Xing Xu, Yue He, Tao Li, Jie Wei, Hongyu Jing, Heqing Wang, Bo Xu, and Hui Song. 2025. A Retrieval-Augmented Generation Framework for Electric Power Industry Question Answering. In *Proceedings of the 2024 2nd International Conference on Electronics, Computers and Communication Technology (CECCT '24)*. Association for Computing Machinery, New York, NY, USA, 95–100. doi:10.1145/3705754.3705771
- [21] Maryamah Maryamah, Muhammad Maula Irfani, Edric Boby Tri Raharjo, Netri Alia Rahmi, Mohammad Ghani, and Indra Kharisma Raharjana. 2024. Chatbots in Academia: A Retrieval-Augmented Generation Approach for Improved Efficient Information Access. In *2024 16th International Conference on Knowledge and Smart Technology (KST)*. 259–264. doi:10.1109/KST61284.2024.10499652
- [22] Laura Moradbakhti, Simon Schreiberlmayr, and Martina Mara. 2022. Do Men Have No Need for “Feminist” Artificial Intelligence? Agentic and Gendered Voice Assistants in the Light of Basic Psychological Needs. *Frontiers in Psychology* 13 (2022). doi:10.3389/fpsyg.2022.855091
- [23] Rodrigo Pereira, Claudio Lima, Tiago Pinto, and Arsénio Reis. 2023. Virtual Assistants in Industry 4.0: A Systematic Literature Review. *Electronics* 12, 19 (2023). doi:10.3390/electronics12194096
- [24] Madeline G. Reinecke, Fransisca Ting, Julian Savulescu, and Ilina Singh. 2025. The Double-Edged Sword of Anthropomorphism in LLMs. *Proceedings* 114, 1 (2025). doi:10.3390/proceedings2025114004
- [25] Jens Reinhardt, Luca Hillen, and Katrin Wolf. 2020. Embedding Conversational Agents into AR: Invisible or with a Realistic Human Body?. In *Proceedings of the Fourteenth International Conference on Tangible, Embedded, and Embodied Interaction* (Sydney NSW, Australia) (TEI '20). Association for Computing Machinery, New York, NY, USA, 299–310. doi:10.1145/3374920.3374956
- [26] Anuschka Schmitt, Thiemo Wambsganss, and Jan Marco Leimeister. 2022. Conversational Agents for Information Retrieval in the Education Domain: A User-Centered Design Investigation. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 486 (Nov. 2022), 22 pages. doi:10.1145/3555587
- [27] Martin Schrepp, Andreas Hinderks, and Jörg Thomaschewski. 2014. Applying the User Experience Questionnaire (UEQ) in Different Evaluation Scenarios. 383–392. doi:10.1007/978-3-319-07668-3_37
- [28] Gayane Sedrakyan, Simone Borsci, Marcos Machado, Patricia Rogetzer, and Martijn Mes. 2025. Design Implications for Integrating AI Chatbot Technology with Learning Management Systems: A Study-based Analysis on Perceived Benefits and Challenges in Higher Education. In *Proceedings of the 2024 International Conference on Artificial Intelligence and Teacher Education (ICAITE '24)*. Association for Computing Machinery, New York, NY, USA, 1–8. doi:10.1145/3702386.3702405
- [29] Francisco I. Teneda, Johnny I. Villacis, Edison G. Espinosa, and Víctor H. Andaluz. 2021. Conversational Agent for Industrial Processes Through Virtual Environments. In *Trends and Applications in Information Systems and Technologies*, Álvaro Rocha, Hojjat Adeli, Gintautas Dzemyda, Fernando Moreira, and Ana Maria Ramalho Correia (Eds.). Springer International Publishing, Cham, 218–229.
- [30] Stefan Wellsandt, Konstantin Klein, Karl Hribernik, Marco Lewandowski, Alexandros Bousdekis, Gregoris Mentzas, and Klaus-Dieter Thoben. 2022. Hybrid-augmented intelligence in predictive maintenance with digital intelligent assistants. *Annual Reviews in Control* 53 (2022), 382–390. doi:10.1016/j.arcontrol.2022.04.001