



Delft University of Technology

Document Version

Final published version

Licence

CC BY

Citation (APA)

Kernan Freire, S., He, T., Wang, C., Niforatos, E., & Bozzon, A. (2026). Factory operators' perspectives on cognitive assistants for knowledge sharing: challenges, risks, and impact on work. *Behaviour and Information Technology*, 45(10), 2248-2275. <https://doi.org/10.1080/0144929X.2025.2574373>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

Factory operators' perspectives on cognitive assistants for knowledge sharing: challenges, risks, and impact on work

S. Kernan Freire, T. He, C. Wang, E. Niforatos & A. Bozzon

To cite this article: S. Kernan Freire, T. He, C. Wang, E. Niforatos & A. Bozzon (2026) Factory operators' perspectives on cognitive assistants for knowledge sharing: challenges, risks, and impact on work, *Behaviour & Information Technology*, 45:10, 2248-2275, DOI: [10.1080/0144929X.2025.2574373](https://doi.org/10.1080/0144929X.2025.2574373)

To link to this article: <https://doi.org/10.1080/0144929X.2025.2574373>



© 2025 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



Published online: 20 Oct 2025.



[Submit your article to this journal](#)



Article views: 1646



[View related articles](#)



[View Crossmark data](#)



Citing articles: 2 [View citing articles](#)

Factory operators' perspectives on cognitive assistants for knowledge sharing: challenges, risks, and impact on work

S. Kernan Freire , T. He , C. Wang , E. Niforatos  and A. Bozzon 

Delft University of Technology, Delft, The Netherlands

ABSTRACT

In the shift towards human-centred manufacturing, our two-year longitudinal study investigates the real-world impact of deploying Cognitive Assistants (CAs) in factories. The CAs were designed to facilitate knowledge sharing among factory operators. Our investigation focussed on smartphone-based voice assistants and LLM-powered chatbots, examining their usability and utility in a real-world factory setting. Based on the qualitative feedback we collected during the deployments of CAs at the factories, we conducted a thematic analysis to investigate the perceptions, challenges, and overall impact on workflow and knowledge sharing. Our results indicate that while CAs have the potential to significantly improve efficiency through knowledge sharing and quicker resolution of production issues, they also introduce concerns around workplace surveillance, the types of knowledge that can be shared, and shortcomings compared to human-to-human knowledge sharing. Additionally, our findings stress the importance of addressing privacy, knowledge contribution burdens, and tensions between factory operators and their managers.

ARTICLE HISTORY

Received 12 May 2025
Accepted 5 October 2025

KEYWORDS

chatbot; manufacturing;
socio-technical factors;
knowledge management;
industry 5.0; human-centred
AI

1. Introduction

An organisation's success hinges on leveraging the knowledge found within the minds of its employees, customers, and suppliers – a process known as Knowledge Management (KM) (Becerra-Fernandez and Sabherwal 2014; Nonaka and Takeuchi 1995). Knowledge management is increasingly recognised as a vital discipline that involves creating, sharing, and using an organisation's knowledge. As such, understanding what leads to a successful Knowledge Management System (KMS) is an important area of research (Becerra-Fernandez and Sabherwal 2014; Chang, Hsu, and Yen 2012).

The importance of KM has motivated efforts to use technology, such as AI assistants, to support human knowledge work, a research field with over 60 years of history. Yet, the success of these systems was often held back by the knowledge acquisition bottleneck (Wagner 2006), capturing and maintaining human knowledge is highly resource intensive. On top of this, social factors such as the trade-off between perceived benefit and effort for the knowledge sharer pose additional challenges (Hoffmann et al. 2022). Domain knowledge can be complex, situational, and ephemeral,

making it difficult to express and store, and it can quickly become outdated (Tan et al. 2006). Yet, cyber-physical systems, such as Cognitive Assistants (CA), show promise in supporting knowledge management (Hoffmann et al. 2022). *In this work, CAs are AI systems that can interact conversationally to support cognitive tasks, for example, in supporting factory operators to share knowledge and learn about problems with the production machines* (Kernan Freire et al. 2022; Leyer, Richter, and Steinhüser 2018).

Due to the complexities of KM, Natural Language Processing (NLP) techniques have had limited use in KM until recently, especially when handling unstructured texts such as manufacturing issue reports (B. Edwards, Zatorsky, and Nayak 2008; Müller, Alexandri, and Metternich 2021; Oruç 2020). Prior work on using assistant systems to share knowledge in factories has focussed on using NLP technologies to deliver instructions efficiently but not the capturing of knowledge. For example, Hoerner, Schamberger, and Bodendorf (2022) presented an assistant system for tacit knowledge sharing in manufacturing but still relies on human experts for the initial knowledge capture. There are examples of gamified experiences to elicit

(manufacturing) knowledge; however, these have not been validated in practice (Balayn et al. 2022; Fenoglio et al. 2022). Therefore, there is a knowledge gap in understanding the real-world challenges and risks when using NLP technologies such as conversational AI for both knowledge capture and sharing on the factory floor.

Recent advances in NLP, such as Large Language Models (LLM) like the GPT-4 (OpenAI 2023) or Claude family,¹ have increased the potential of CA systems. A foundational LLM can offer various functions, such as answering general knowledge questions, refining text, and checking logic. Additionally, the information contained in the foundational models can be extended by providing it with context material, a process called Retrieval Augmented Generation (RAG) (Lewis et al. 2020). LLMs are capable of more sophisticated, flexible, and human-like reasoning than previously possible with AI. These attributes could help overcome issues with agential AI during work support, such as not considering the social and context aspects of knowledge sharing (Cabitza, Campagner, and Simone 2021; Shneiderman 2022).

Operating a modern production line is knowledge-intensive, dynamic, and socio-technical in nature (Kotthaus et al. 2022). As such, it is important to follow a human-centric approach when deploying new AI tools (Shneiderman 2022). Work by Hoffmann et al. (2022) and Mörike (2022), has shown how the social and hierarchical structures in manufacturing play a key role in knowledge sharing practices. Yet, one of the challenges of conducting research in manufacturing is the restrictions management imposes due to safety or a need for oversight, such as limited time to interview operators or requiring that a manager always be present. Despite the potential to influence how open the study participants are in discussing work matters, this is often unavoidable as factory management is vital to coordinating any research in factories (Cheon, Schneiders, and Skov 2022). Thus, the (unfiltered) perspective of factory operators is poorly represented in literature.

In tackling these gaps, this work explores the human, organisational, and practical factors surrounding deploying CA systems in the real-world factory context from both the operator's and managers' perspectives, leading to the following research question:

What are factory operators' and management perceptions of the impact and socio-technical risks and challenges of using cognitive assistants for knowledge sharing?

The HCI and CSCW fields have long contributed to Knowledge Management, taking a socio-technical and practice-based perspective (Ackerman et al. 2013). We aim to expand on this body of knowledge by examining

the latest generation of AI knowledge sharing tools for manufacturing, (LLM-powered) CAs. Over a period of two years, we deployed CAs in four phases, ranging from technology probes simulated by researchers to voice assistants connected to a live production line to LLM-powered chatbots. Using feedback from system evaluations and responses from semi-structured interviews conducted at two factories, we present findings from a hybrid deductive/inductive thematic analysis of $N=40$ operators' and management's perceptions of CAs.

Through our longitudinal study, we aim to contribute to a better understanding of the challenges and impact of using CAs in manufacturing. Thus supporting future research and implementation of CAs for knowledge sharing that are both effective, and sensitive to the socio-technical context of operating complex systems in factories. Our contributions can be summarised as follows:

- (1) We provide a deeper understanding of the opportunities, challenges, and risks associated with using (LLM-powered) CAs by factory operators from both operator and management perspectives.
- (2) We provide insights on the tensions between factory operators and management regarding technology-facilitated knowledge sharing.
- (3) We present design guidelines for effective (LLM-powered) CAs in manufacturing settings.

2. Background and related work

This section provides an overview of current research in Knowledge Management (KM) and the utilisation of technology, specifically Cognitive Assistants (CA), to enhance knowledge sharing in organisational settings. It discusses the challenges and advancements in KM practices, the significant role of technological interventions in facilitating knowledge acquisition, dissemination, and application, and the socio-technical aspects that influence implementing these systems in manufacturing environments.

2.1. Knowledge management

Organizational KM involves leveraging the knowledge of employees, customers, and suppliers (Becerra-Fernandez and Sabherwal 2014). This entails processes to acquire, manage, and share knowledge within an organisation to achieve a competitive advantage (Nonaka and Takeuchi 1995). In practice, the core tenet of KM is to 'save content, it will be useful later' (Abubakar et al. 2019). For this work, CAs were deployed to support

knowledge sharing between factory operators. In the following sections, we will outline the key processes surrounding knowledge sharing when facilitated by a digital system such as a CA, namely, knowledge acquisition, sharing, and application.

Knowledge acquisition involves eliciting knowledge, explicating it, and formalising it for later use (Cooke 1994). Knowledge can be elicited once it has been created in the mind of an employee. *Eliciting* involves extracting knowledge from domain experts or sources, often through interviews, surveys, observation, or analysis of existing documentation. Elicitation can uncover both explicit and tacit knowledge, understanding not just what is known but also how it's applied in practice. In this work, we focus on explicit knowledge that can be readily expressed in words as opposed to tacit knowledge that can not (Nonaka 2009).

Once knowledge has been elicited, it needs to be made explicit. *Explication* involves articulating the knowledge clearly and understandably. This may involve defining terms and organising information into coherent structures. Explication helps ensure that knowledge is comprehensible to others and can be effectively communicated.

Finally, during *formalization*, knowledge is structured and represented in a formal language or framework. Formalization makes knowledge more precise and facilitates analysis and manipulation by digital systems. This might involve creating models, rules, ontologies, or other formal knowledge representations. Formalization helps facilitate (AI) reasoning and decision-making based on the captured knowledge.

Knowledge sharing is defined as the process of transferring knowledge between individuals, groups, or organisations (Alavi and Leidner 2001). Once knowledge has been acquired, it can be shared with others who can benefit from it. Traditionally, this can be done through various channels such as documentation, training sessions, workshops, seminars, meetings, or digital platforms. While many of these strategies aim to stimulate face-to-face knowledge sharing, organisations are also interested in systems that formalise the knowledge. This ensures that the knowledge persists beyond employees' memory and facilitates sharing at scale and asynchronously. In this work, we explore using CAs that act as an intermediary to help efficiently share newly created knowledge.

Knowledge application is when knowledge is used to solve problems and make informed decisions to profit the organisation (Probst, Raub, and Romhardt 2000). This involves analyzing the situation, identifying relevant knowledge, and applying it effectively to address the issue at hand. In the context of this work, the CAs contribute

primarily to the process of identifying and retrieving relevant knowledge. Analyzing the situation and taking action is still largely up to the human operator.

Knowledge management systems typically face many human, organisational, and technical challenges, such as employees hiding their knowledge to maintain (perceived) power or inefficient knowledge sharing practices (Becerra-Fernandez and Sabherwal 2014). Addressing these challenges, organisations have developed their own methods and systems to facilitate knowledge sharing.

2.2. Knowledge sharing systems

Designing KS systems requires consideration of a broad range of factors, often intertwined with social and human elements. Organizations across various sectors have developed their own dedicated KS systems. For instance, the 'Inside IBM' project used AI, information systems, and user-centred design to aggregate IBM's accumulated product support knowledge into a single system (Massey, Montoya-Weiss, and Holcom 2001). Siemens employs TechnoWeb for social collaboration, which leverages a combination of social networking tools, collaborative platforms, and knowledge repositories to facilitate the sharing and management of expertise across the organisation (Lakhani et al. 2013). TechnoWeb integrates features such as discussion forums, expert directories, and document sharing to create a dynamic and interactive environment for knowledge exchange. NASA uses the Lessons Learned Information System (LLIS)² to document and share project experiences. LLIS employs a structured database approach, where lessons are systematically categorised, indexed, and made searchable to ensure that critical insights from past projects are easily accessible for future missions (Liebowitz 2002). This system includes detailed metadata tagging, cross-referencing of related lessons, and a review process to validate the accuracy and relevance of the information. Furthermore, the enhanced version of LLIS that the Goddard Space Flight Center uses employs an 'active' push feature to recommend lessons that match the knowledge needs of specific individuals. *These diverse systems illustrate the various methods employed to manage and leverage knowledge within organisations, ranging from searchable lessons learned and social collaboration tools to active knowledge recommendations and systematic documentation processes.*

Several factors contribute to the success of KS systems, as exemplified by the cases of Google, Toyota, and Xerox's Eureka project described below. In the case of Google, they emphasise psychological safety

for KS (Duhigg 2016). Toyota, perhaps KM's most renowned manufacturing success story, employs network-level knowledge-sharing processes. These processes effectively engaged members in sharing critical knowledge, deterring free-riding behaviour, and minimising the obstacles to locating and obtaining valuable knowledge (Dyer and Nobeoka 2000). The Eureka project at Xerox showcases the successful use of technology to gather, validate, and share best practices across their customer service engineers (J. S. Brown and Duguid 2000). By involving customer service engineers throughout the design process and encouraging local champions to support adoption and feedback, Xerox contributed to the project's success (Bobrow and Whalen 2002). Ultimately, the Eureka project is said to have saved Xerox \$100 million (Bobrow and Whalen 2002; J. S. Brown and Duguid 2000). *Overall, these cases highlight the importance of minimising the burden on knowledge contributors, ensuring high knowledge quality, and involving end-users throughout the design process to address their needs.*

2.3. Integrating AI in knowledge sharing systems

While Eureka was an early example of an online knowledge base that service engineers could manually record and retrieve reports about machine problems, expert systems were the first widely used AI systems to reason over a knowledge base to support human decision-making. Using AI assistants for KM has progressed significantly from early expert systems in the 1970s, such as Dendral (Lindsay et al. 1993) and MYCIN (Shortliffe et al. 1975), which were designed to answer questions in specific domains based on rules defined by domain experts. In recent years, the improvements in natural language processing and deep learning have given rise to more sophisticated AI assistants, such as IBM's Watson (High 2012).

KS systems incorporating AI and Internet of Things (IoT) technologies are transforming KS within manufacturing environments (Casillo et al. 2020). Yet, despite the promise of these technologies to enhance KS systems, several socio-technical hurdles persist. For instance, significant resources are required to develop and maintain knowledge bases (Cullen and Bryman 1988), and data quality issues often undermine efforts to automatically uncover knowledge from existing unstructured issue reports (B. Edwards, Zatorsky, and Nayak 2008; Müller, Alexandi, and Metternich 2021; Oruç 2020). Similar challenges have faced the use of NLP in healthcare, such as IBM Watson's Oncology Expert Advisor in 2016, which searched through 'unstructured' physician notes to help suggest treatments. Yet, Watson's NLP failed to effectively process

the jargon, shorthand, and subjective comments or pick up the nuances that human physicians could, demonstrating the challenges in handling human-generated texts (Strickland 2019). These challenges are socio-technical because they involve both technical aspects, such as data quality and the capabilities of NLP in processing it, and human aspects, such as the effort put toward documentation, which inhibits the data's utility for KS (Sexton et al. 2017). *With recent advancements in NLP, such as the advent of LLMs such as GPT-4 (OpenAI 2023), it may be possible to overcome some of the hurdles of leveraging unstructured human documentation faced by older systems. Additionally, in tackling the human aspect, there is a knowledge gap in designing KS systems that facilitate the acquisition of human knowledge while balancing the (perceived) effort of knowledge authoring such that factory operators are consistently engaged in contributing useful knowledge.*

Foundational LLMs do not contain context-specific knowledge, for example, how to fix a machine in a specific factory, they possess extensive general knowledge and reasoning abilities (Zhao et al. 2023), enabling them to excel in processing complex information (Jawahar, Sagot, and Seddah 2019), generate insights and reasoning (Wei, Tay, et al. 2022). While an LLM can be trained from the ground up for a specific domain, such as medicine, this is typically cost-prohibitive for individual organisations. Even so, using a domain-specific LLM would have two key challenges: (1) reliance on outdated information from their pre-training data, and (2) potential inaccuracies in factual content, a phenomenon termed 'hallucination' (Bang et al. 2023; Zhao et al. 2023). To mitigate these issues and maximise the utility of LLMs in specialised, knowledge-intensive tasks, techniques such as fine-tuning, chain-of-thought (Wei, Wang, et al. 2022) few-shot prompting (Gao, Fisch, and Chen 2021; T. Brown et al. 2020), and Retrieval Augmented Generation (RAG) (Lewis et al. 2020) can be adopted. RAG involves retrieving relevant information from a knowledge base before generating an appropriate response for a human. RAG is an efficient way to harness the advanced NLP and reasoning abilities of LLMs to answer people's questions on domain-specific topics without needing to train a new model from scratch. Furthermore, the availability of source material, the information retrieved from the knowledge base, improves transparency and enables fact-checking, key aspects in fostering trust and acceptance in AI systems (Vorm and Combs n.d.). *As such, foundational LLMs used in conjunction with transparent knowledge bases hold significant potential to enhance organisational KM.*

2.4. Cognitive assistants for factory operators

In recent years, research into AI assistants in manufacturing has increased, taking advantage of recent advancements in NLP and the proliferation of IoT to support decision-making, operational efficiency, and training (Casillo et al. 2020; Castane et al. 2023; Colabianchi, Tedeschi, and Costantino 2023; Hannola et al. 2020; Hoerner, Schamberger, and Bodendorf 2022; Kiangala and Wang 2024; Ras et al. 2017; Tao et al. 2019)). Although the social and human aspects are largely missing (Hannola et al. 2020), they have been receiving increasing attention as demonstrated by the recent paradigm of Industry 5.0 where human well-being is key (Breque, de Nul, and Petridis 2021; Xu et al. 2021). In the following sections, we discuss research on AI assistants in manufacturing, the use of knowledge bases, knowledge acquisition from humans for knowledge sharing, and the associated socio-technical challenges.

2.5. AI assistants

In factories, human operators collaborating with AI assistants can be viewed as a socio-technical system consisting of the operator, the operator's tasks in the factory, and the technology enabling the assistant (Maedche et al. 2019). AI assistants can help operators fix problems with the machines or set them up optimally by enabling efficient information retrieval, remote machine control, decision support, and sharing knowledge (Berger and von Garrel 2023; Castane et al. 2023; Dhuieb, Laroche, and Bernard 2013; Li et al. 2022). To enhance their capabilities, AI assistants can be integrated into factory systems, such as control and monitoring systems for production machines, scheduling systems, and sensors. Some systems make use of the available data to perform predictive maintenance using Machine Learning (ML) (Wellsandt et al. 2022) or NLP to extract and represent knowledge from texts (Naqvi et al. 2022). To interact with operators, most AI assistants use conversational AI, a chat or voice interface, to provide efficient and natural ways for operators to communicate.

Recently, manufacturing companies have cautiously adopted advanced AI assistants that use LLMs. For instance, Mercedes-Benz (2023) integrated ChatGPT into vehicle production to improve error identification, quality management, and process optimisation. This allowed quality engineers to simplify complex evaluations and presentations of production data through dialogue-based queries. Furthermore, Xia et al. (2023) showcased how in-context learning and the injection

of task-specific knowledge into LLMs can improve the planning and control of production processes. *However, systems like these rely on an up-to-date and detailed knowledge base, a major hurdle for any KMS.*

2.6. Knowledge-intensive AI assistants

The knowledge bases and rules that inform AI assistants in manufacturing are usually static, having been defined by a domain expert during development (Casillo et al. 2020; Rooein et al. 2020; Trappey et al. 2022; Wieland et al. 2016). Some systems use live data but are focussed on presenting the status of production systems and perform (limited) reasoning over the data (Baldauf et al. 2021; Li and Yang 2021; Mantravadi, Jansson, and Moller 2020; Pavlov et al. 2020; Reis et al. 2022). For example, Casillo et al. (2020) presents a chatbot for training new operators based on a predefined curriculum, obtaining promising results. Baldauf et al. (2021) deployed smartwatches that notified operators of machine errors, citing the importance of concise information display. Additionally, the chatbot demonstrated by Trappey et al. (2022) utilises VR to deliver knowledge in response to FAQs (frequently asked questions). *Overall, existing research on AI assistants focussed on understanding how to deliver knowledge or present (live) information to operators but not acquire knowledge from operators.*

2.7. Cognitive assistants: acquiring and sharing human knowledge

In the context of this study, *Cognitive Assistants (CAs) are an advanced type of AI assistant that interacts conversationally to acquire knowledge from humans and share it with other humans, supporting decision-making.* Cognition refers to the mental processes of acquiring and comprehending knowledge (Angulo, Chacón, and Ponsa 2023), which the CA aims to support. To do so effectively, a CA can use several technologies, including dialogue management, NLP, context awareness, databases, and ontologies. While deploying CAs at a production line appears novel, we discuss notable works that evaluated (parts of) the underlying technologies and techniques in the following sections.

Acquiring knowledge using conversational AI and/or mobile devices has been explored and shown to be promising. For example, Fenoglio et al. (2022) introduced a role-playing game involving virtual agents, human experts, and knowledge engineers to refine knowledge graphs that were algorithmically generated. The authors raised several important ethical concerns regarding the processing of personally identifying

data, deciding to avoid audio and video recordings, and other concerns regarding misuse of employee monitoring. Hoerner, Schamberger, and Bodendorf (2022) built a digital assistance system to support operator troubleshooting processes on the shop floor using captured knowledge. While they allowed operators to suggest edits, *knowledge engineers still performed the initial knowledge capture process*. In contrast, Hannola et al. (2020) deployed smartphone applications for factory operators to take notes and videos of solutions to production line problems themselves, thus enabling other operators to access newly created knowledge. *While this work highlighted the potential of facilitating KS through digital technologies, it did not investigate the application of LLMs and/or conversational AI*.

Examples outside of manufacturing demonstrated the successful use of conversational AI to crowdsource knowledge acquisition, namely, a game to elicit knowledge from crowdworkers (Balayn et al. 2022) and a context-aware chatbot that simultaneously fulfils information needs while acquiring new knowledge (Bradeško et al. 2017). A key consideration for these systems is ensuring the quality of acquired knowledge, thus employing rating or validation mechanisms. *Although these conversational AI systems were not deployed in a manufacturing environment, they demonstrate the potential of using conversational interactions to efficiently acquire knowledge*.

2.8. Socio-technical challenges

While many KMS have been deployed in manufacturing, the human potential for knowledge sharing and the associated social factors have been largely overlooked (Hannola et al. 2020). Work in factories is socio-technical rather than purely technical, deterministic activities (Kotthaus et al. 2022; Mucha et al. 2018). While it is clear that digital tools can facilitate knowledge sharing among users, for example, by enabling factory operators to post comments under existing work instructions to facilitate discussion about best practices (Ludwig, Boden, and Pipek 2017), these tools can have many unforeseen social effects and barriers. For example, introducing relatively simple tools, such as a machine repair ticketing system, can shift the balance of control towards or away from operators depending on who is given privileges to create, view, and modify tickets (Kotthaus et al. 2022). Furthermore, operators omitted their valuable informal knowledge when asked to create learning material collaboratively in favour of the standard working procedures (Weinert et al. 2022). Additionally, requiring operators to use smartphones to document their work can be perceived

as controlling (Hannola et al. 2020). Thus, these factors emphasise the importance of considering social factors when introducing new digital knowledge sharing practices, which are generally overlooked when designing AI assistant systems (Cabitza, Campagner, and Simone 2021).

To tackle the complex social challenges discussed above, it is crucial to look beyond the technical aspects when developing AI tools for knowledge sharing. Although not social in nature, considering human-computer interaction factors such as avoiding information overload and ensuring the information needs are met on an individual level are also key (Hoffmann et al. 2022). Indeed, recent work on assistance systems for sharing tacit knowledge in factories that had focussed on technological aspects emphasised the importance of comprehensively evaluating the social, psychological, and organisational implications (Hoerner, Schamberger, and Bodendorf 2022). The crucial importance of socio-technical factors is what motivated us to holistically reflect on the feedback we received from operators and managers while deploying CAs, culminating in this study.

3. Real-world case studies in factories

The factories manufactured detergents on production lines that could rapidly switch between over a hundred detergent types depending on customer demand. The factories operated 24/7 with three eight-hour shifts per day, and each production line was typically manned by two operators. The operators reconfigure, operate, and resolve issues with the production lines, all knowledge-intensive tasks, especially considering the frequent production and quality issues.

The development of the CAs used in this study was initiated according to objectives set by the factory management: to increase production performance and reduce training time by supporting (tacit) knowledge sharing amongst operators and training operators to follow standardised procedures. Despite the top-down initiation of the project, we involved the operators throughout the development and evaluation of the system to ensure that we were also meeting their needs and values. This research was approved by the factory operators' councils and our institution's human research ethics committee.

3.1. Artefacts

We evaluated four CAs of varying levels of sophistication and functionality, from technology probes simulated by researchers to CA systems connected to a live

Table 1. CA capabilities evaluated per phase.

No.	Capability	Phase 1	Phase 2	Phase 3	Phase 4
1	Speech input	✓		✓	✓
2	Speech output	✓		✓	
3	Text input		✓	✓	✓
4	Text output		✓	✓	✓
5	Capture issue handling knowledge	✓	✓	✓	✓
6	Share issue handling knowledge		✓	✓	✓
7	Automatic validation of captured knowledge				✓
8	Human approval of captured knowledge				✓
9	Capture product knowledge		✓	✓	
10	Rate captured product knowledge		✓	✓	
11	Approve shared product knowledge		✓	✓	
12	Share product knowledge		✓	✓	
13	Graph machine data for systematic reflection			✓	
14	Machine vision for context-aware dialogue		✓	✓	
15	Upload new documents to the knowledge base				✓
16	Answer predefined FAQs		✓	✓	
17	Provide relevant standard work instructions		✓	✓	
18	Answer queries using factory documentation				✓
19	Real-time explainable machine settings advice			✓	

production line, allowing us to explore many facets of the operators' experiences. The evaluated functions are listed in Table 1.

3.1.1. Phase one: technology probe during production line operations

During the first phase, we used a technology probe (Boehner et al. 2007) to evaluate the user experience of using a CA for capturing knowledge at the production line. The CA was simulated by a researcher, with operators interacting verbally via a Bluetooth headset (see Figure 1). The protocol followed a form-filling conversation to capture issue-handling knowledge.

3.1.2. Phase two: CA with a simulated production line

For phase two, we developed a CA that used live data from a simulated production line and human location tracking using stereoscopic cameras for context awareness. The CA was text-based, built using RasaX, and connected to a Neo4j knowledge graph. The knowledge base was prepopulated with a 100 issue reports captured during phase 1. The interface facilitated efficient user interactions by prefilling responses as buttons based on proximity to machines and live data (see Figure 2(a)).



Figure 1. Headset used at the production line for data collection during phase 1.

3.1.3. Phase three: CA with partial LLM integration connected to the production line

In phase three, we evaluated a CA capable of both voice and text interaction, with access to live machine data. New features included ML-based recommendations, production statistics, and graphical representations. The most significant change was the integration of an LLM-powered Retrieval Augmented Generation (RAG) function using the GPT-3.5 API, accessible via an Android application.

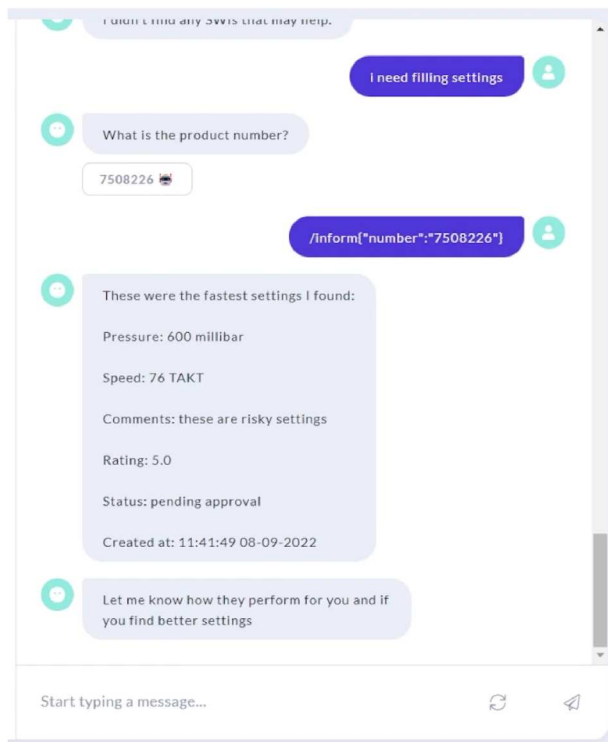
3.1.4. Phase four: LLM-powered cognitive assistant

The final phase focussed on LLM-powered functions, using the gpt-3.5-turbo-0613 model API, LlamaIndex for response generation, and Gradio for the browser-based frontend (see Figure 3(b)). New functions included RAG for user queries, digital forms for issue analysis reports, automatic and human validation of reports, and document upload capabilities.


3.2. Research approach

This study adopted a *longitudinal, action research approach* to investigate the development, deployment, and use of conversational agents (CAs) in two European detergent factories over a two-year period. While there is no cohesive definition or shared understanding of what *longitudinal* entails in HCI research (Kjærup et al. 2021), we use it here to describe a multiphase study during which we collaboratively designed artefacts, deployed and tested them over two years at two factories.

The longitudinal nature of the research allowed us to reduce the impact of the novelty effect, capture changes over time, elicit honest feedback having built trust with the participants, and explore the impact of new



(a)



(b)

Figure 2. (a) The Rasa X user interface from phase 2 and (b) A simulated 'weight checker' user interface that was connected to the assistant during phase 2. Each bar represents the weight of a canister that was filled with detergent. The 'weight checker' weighs each canister and stops the production line if the weight falls below a specified threshold.

technologies. By conducting multiple evaluation phases with increasingly sophisticated CA prototypes, we were able to iteratively refine both the artefacts and our understanding of the sociotechnical context. This extended engagement provided insights into the adoption process, barriers and enablers, and the relationship between technology and work practices.

Action research was chosen as it emphasises iterative cycles of planning, action, observation, and reflection, with the dual aims of generating practical solutions and advancing scientific understanding (Baskerville 1999). This approach enabled close participatory collaboration with both factory management and operators, ensuring that the evolving CA systems addressed real-world needs and values as they emerged in practice (Greenbaum and Kyng 1991). Furthermore, this collaborative approach aimed to improve the relevance and acceptance of the technology while also improving our understanding of the socio-technical context.

We adopted a flexible and adaptive methodology depending on the phase of the project. This flexibility helped adapt to the dynamic environment of 24/7 production lines, the diversity of operator expertise,

evolving capabilities of the technology, and the shifting priorities of factory management.

Ethical approval for the research was obtained from the factory operators' councils and our institution's human research ethics committee.

3.3. Participants

We conducted seven evaluations with prototypes at two factories with a total of $N = 40$ participants, of which $n = 5$ were female and $n = 36$ were male. The participants included operators, managers, and technicians, as detailed in Table 2 and in the list organised by phase below:

During *phase 1*, six ($n = 6$) male operators participated, with occasional observations by two managers and a technician. The six operators, comprising 3 shifts of 2 operators, were recruited in person on the days of the experiment after consultation with their shift lead and production manager. Each shift typically comprises an experienced, lead operator paired with a less experienced operator.

For *phase 2*, three focus group sessions included 11 ($n = 11$) factory operators and managers (2 female, 9 male).

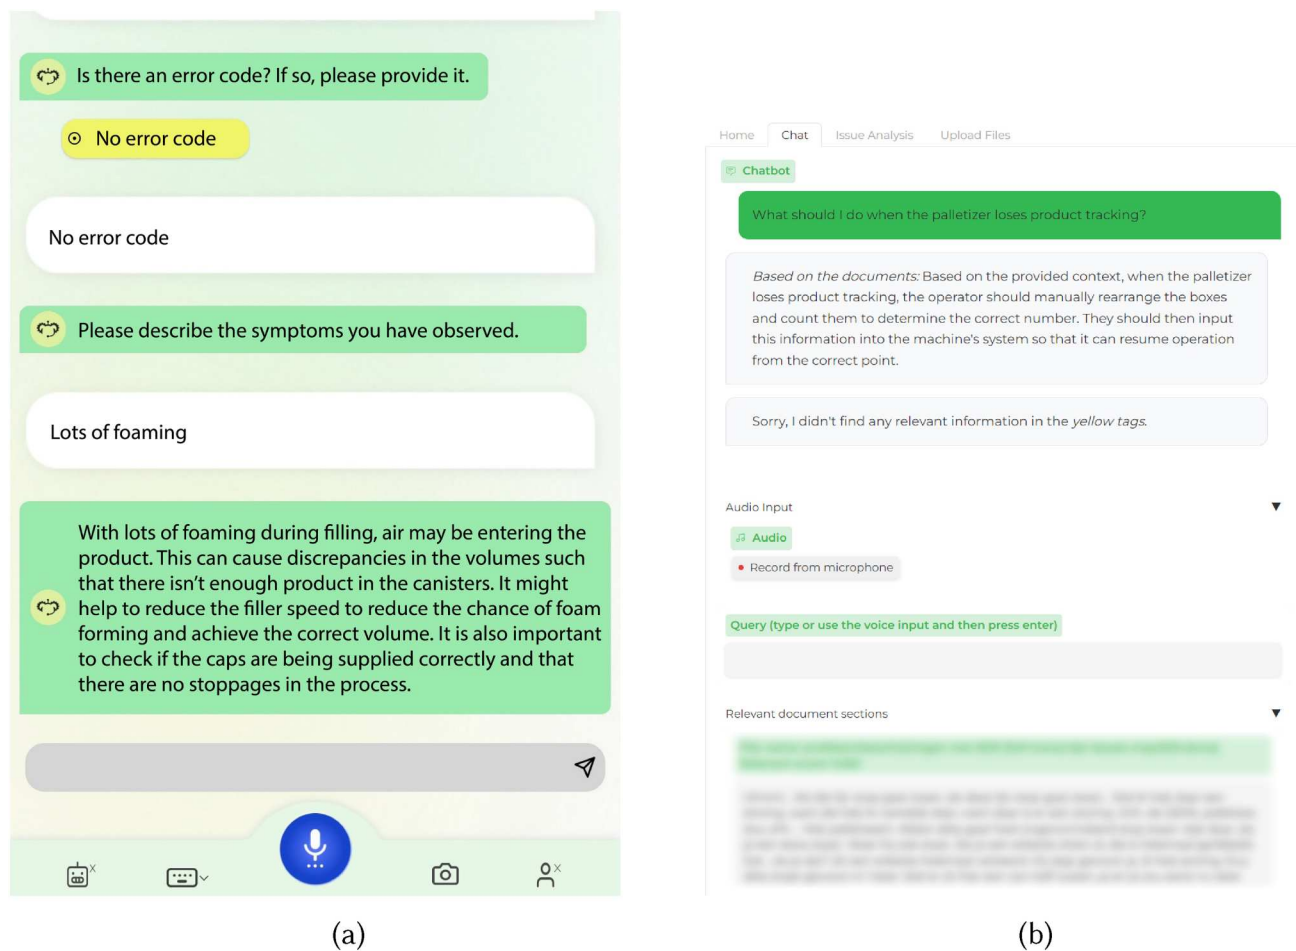


Figure 3. (a) Android user interface for phase 3 and 4 and (b) The query tab of the LLM-powered CA from phase 4.

In *phase 3*, five ($n = 5$) operators and three ($n = 3$) managers, all male.

Finally, during *phase four*, three ($n = 3$) operators and one ($n = 1$) manager from the first factory, and four ($n = 4$) managers and five ($n = 5$) operators from the second factory.

3.4. Procedure and data collection

In preparation for *phase 1* (Jan. 2021–May 2021), operators and managers described current knowledge management practices. Over three days, operators shared how they solved issues in realtime with a CA being simulated by a researcher. Participants used a Bluetooth headset to interact verbally whilst at the production line. A researcher simulated the CA using a protocol representing state-of-the-art NLU and dialogue management. Data was collected over three days. Researchers recorded observations by managers and operator feedback during interactions with the CA, resulting in 18 ($n_c = 18$) comments relevant to using CAs in factories.

For *phase 2*, we used a browser based chat interface in addition to a computational model to simulate the production line system. Stereoscopic cameras were installed at the factory for context awareness for the CA. Three focus group sessions were conducted, each including a video presentation, a structured user test with eight tasks, and a semi-structured user test. Participants wrote post-it notes providing feedback on what worked well, what didn't and concrete suggestions for improvements. After individually writing messages on postits, researchers also moderated group discussions to unpack some of the key topics and elicit further feedback. All the feedback was collected as anonymous post-it notes, resulting in 55 ($n_c = 55$) comments relevant to this study.

During *phase 3*, the new mobile-based CA was demonstrated to operators and managers. Then, operators used the CA at the production line as they saw fit for 15–30 minutes, while managers completed five tasks in an office setting. Tasks included saving and retrieving machine parameters, generating production statistics, and using the LLM-powered decision support for solving problems. Upon completion of the

Table 2. Participant Data.

No.	Role	Factory	Phase	Test duration (min)	Prior Phases	Gender
1	Operator	1	1	60–180	No	Male
2	Manager	1	1	0	No	Male
3	Operator	1	1	60–180	No	Male
4	Manager	1	1	0	No	Male
5	Operator	1	1	60–180	No	Male
6	Operator	1	1	60–180	No	Male
7	Technician	1	1	0	No	Male
8*	Mixed	1	2	10	No	Mixed
21	Manager	1	3	15–30	Phase 2	Male
22	Operator	1	3	15–30	No	Male
23	Technician	1	3	15–30	No	Male
24	Operator	1	3	15–30	No	Male
25	Operator	1	3	15–30	No	Male
26	Manager	1	3	15–30	Phase 2	Male
27	Manager	2	4	15–30	No	Female
28	Manager	2	4	15–30	No	Female
29	Manager	2	4	15–30	No	Male
30	Manager	2	4	15–30	No	Female
31	Operator	1	4	15–30	No	Male
32	Operator	1	4	15–30	No	Male
33	Manager	1	4	15–30	Phase 2,3	Male
34	Operator	1	4	15–30	No	Male
35	Manager	1	4	15–30	Phase 1,2,3	Male
36	Operator	2	4	15–30	No	Male
37	Operator	2	4	15–30	No	Male
38	Operator	2	4	15–30	No	Male
39	Operator	2	4	15–30	Phase 3	Male
40	Operator	2	4	15–30	Phase 1	Male

experimentation and tests, we conducted semi-structured interviews on their experiences with the CA, eliciting feedback, the potential impact of the CA, and ideas for improvements. Relevant statements were extracted from interviews, resulting in 60 ($n_c = 60$) comments.

For *phase 4*, evaluations were conducted both in person (first factory) and asynchronously (second factory). Participants were first assigned tasks to evaluate knowledge capture and sharing capabilities of the CA, such as asking for help with recent problems and recording solutions. Operators were also given the opportunity to experiment with the CA as they saw fit while at the production line. After completing the tasks and experimentation in factory one, researchers conducted a semi-structured interview on the perceived benefits, risks, adoption barriers, operator-management relations, and suggestions for improvements. For factory two, the participants provided their responses in a survey. The interviews and surveys resulted in 118 ($n_c = 118$) comments that were relevant for this study.

3.5. Data analysis

The analysis was conducted based on 251 ($n_c = 251$) comments extracted from interviews, survey answers, and post-it notes collected during phases one through four. All comments related to the user experience, social,

technical, and organisational factors of the CA deployments at the factories were included. Unintelligible comments were excluded. Complex statements comprising multiple standalone arguments or observations were divided into separate comments to facilitate sorting and analysis. If other parts of the original statement provided relevant context, this was included in regular brackets.

After several rounds of qualitative data collection over two years, we conducted a hybrid deductive/inductive thematic analysis with two independent coders as described by Fereday and Muir-Cochrane (2006) and recently applied by Martinez-Maldonado et al. (2023). The coders consisted of two experts in HCI, one of whom was integrally involved in the entire project, and second independent coder who was familiar with the project. The high-level themes were defined deductively while the subthemes were defined inductively from the corpus.

The first deductive stage of the thematic analysis involved defining themes stemming from the research question and literature review, resulting in six themes: (1) Impact on Work Experience, (2) Optimizing Knowledge Sharing, (3) Adoption and Change Management, (4) Privacy, Safety, and Ethics, (5) Usability and User Experience, and (6) Technical and Operational Issues. The two coders worked independently to code all the comments under the 6 main themes, iterating on their definitions and resolving disagreements across several sessions.

Once the coders achieved consensus on the distribution of the comments among the 6 main themes, they moved onto the inductive phase of the analysis to identify subthemes. Again, the two coders worked independently, analysing all comments within each theme to identify subthemes. After independent coding, the coders met to discuss their findings, merging and adjusting subthemes as deemed necessary until a consensus was reached. This process resulted in 20 subthemes across the 6 main themes.

While some comments were brief, the corpus also contained rich descriptive accounts of specific incidents and insights. The mean word count for a single comment was 16 ($SD = 9.56$), ranging from a minimum 2 words to a maximum 65 words. In part to accommodate the richer, complex comments, we agreed to allow comments to be sorted into multiple themes and subthemes if applicable.

4. Results

Our results are organised by themes to answer the research question: ‘What are factory operators’ and management perceptions of the impact and socio-

technical risks and challenges of using cognitive assistants (CAs) for knowledge sharing?”. Each subsection below corresponds to the themes and subthemes that arose from the thematic analysis. Together, these sections address elements in the research question on the perceived impact, risks, and challenges, viewed from both operator and management perspectives. A mapping table below depicts what part of the research questions are addressed by which (sub)themes. Theme 4.1 speak to impact, themes 4.2–4.6 to socio-technical risks and challenges, and subtheme 4.3.3 on the operator and management perspective. The operator and management perspective is also addressed across several other (sub)themes where applicable.

4.1. Impact on work experience

This section addresses the research question by exploring how both factory operators and management perceive the impact of cognitive assistants on daily work experience, including information retrieval, problem-solving, knowledge sharing, and training.

4.1.1. Improved information retrieval

The comments suggest that the system positively impacts information retrieval within the organisation. The CA was perceived to provide immediate answers to doubts about machine operation, making access to company documents more user-friendly (P30), and improving the flow of information (P28). Furthermore, participants saw it as a tool that modernises the factory (P27). However, one concern is that operators may benefit more from traditional information retrieval systems (P40).

4.1.2. Useful for problem-solving

Many participants ($n = 9$) are positive about the usefulness of the system and specific functions (P22, P24, P25, P36, P30, P31, P32, P33, and P34), and some participants, such as P25, make a distinction between some functions being useful and others not. Specifically, the system provides greater speed in carrying out some small tasks (P29), it has a well-set-up logical control function (P30), and helps solving problems quickly (P8*, P24, P25, and P37) as stated here ‘I find it nice that you can look back. What did we do last time with the same problem? That, I think, is a significant advantage’ by P24, and ‘[the CA] can be a great tool in problem solving in the day-by-day activities’ by P8*. However, there is a disadvantage of wasting time looking for a solution to a problem if it is not reported in the system’s history (P29).

4.1.3. Effective knowledge sharing

Several benefits regarding knowledge sharing are mentioned by the participants, namely that it is useful to transfer knowledge from experienced to novice operators (P31), ‘The system also allows for the communication of tacit knowledge between teams’ (P33), it can also help individuals recall knowledge of how they previously solved issues (P24, P32) and experts’ knowledge can now persist and is always available (P34). These comments show that participants generally see it as a mechanism for experts to share their knowledge with novices and as a memory augmentation for themselves.

That said, there are concerns that receiving information from the system could be less quick, effective, and comprehensive than asking the expert staff present (P27). This is likely true, assuming that there is an expert present with the relevant knowledge. However, our observations at the production line and statements by operators, such as ‘A few hours are lost every week on previously solved issues’ (P33) and P34 stated that operators ‘Lose about half an hour per day over one team per line’ suggests that information on previously solved issues was not readily available.

Whereas management believes in standardising approaches so every operator operates at the highest achievable performance, some operators do not think this is feasible. One operator in particular, P25, was skeptical of the benefits of knowledge sharing as stated here, ‘We always adjust to the specific product characteristics and line conditions’, ‘We often have different issues’ and ‘Because everyone has their own approach, some run slower, and others have their own techniques’. In other words, their work is so complex that the ‘best practice’ is highly dynamic and dependent on context factors and individual strategies. This reflects experienced operators’ pride in the strategies they developed, which may hinder the effectiveness of knowledge sharing. Conversely, P22 mentioned that ‘It would be useful to know if and how other operators have produced at higher speeds’ (P22), suggesting they would be interested in learning from each other. This reflects a split among operators regarding the perceived usefulness of knowledge sharing.

4.1.4. Training time

Opinions are divided on whether training time will be reduced, with P32 stating that ‘[the CA] will address staff shortages and training time’ whereas P31 stated that training will not directly be reduced but independence will be increased. Others, such as P33, were unsure whether training time would be impacted but expect improvements as the technology matures.

Generally, participants thought that novice operators would benefit most (P8*, P22, P24, P25, P31, P34, and P38); for example, P22 stated, 'This is very nice, yes, especially for beginner [operators]' and P25 stated, 'It's useful for newcomers who might not have as much product and line knowledge'. These statements support the idea that independence can be improved and training time can be reduced.

4.2. Optimizing knowledge sharing

This section addresses the socio-technical challenges of sharing knowledge via cognitive assistants. Furthermore, this section starts with outlining the need for efficient knowledge sharing for operators on a production line.

4.2.1. Value of efficient knowledge sharing

The comments highlight the importance of expert operator knowledge at work. This know-how is not captured in manuals or documentation; for example, 'An operator poured water over one of the canisters in the filling machine when the line had stopped to get it moving again' (P6) and 'We adapt to the product's characteristics, and some products require slower operation due to foam production' (P25). Yet, existing mechanisms for capturing this knowledge were largely unused, as stated by P2: 'Because the problems are so poorly documented in [current issue reporting system], we don't have a good overview of the problems' P25 also describes how they develop their own strategies: 'Because everyone has their own approach, some run slower, and others have their own techniques' and the necessity to adapt to changing context: 'Yes, sometimes we have products that change their speed, or the machine is modified, allowing faster operation'. This partially explains why existing tools for documenting knowledge are quickly outdated and resource-intensive to maintain, as discussed in Section 4.2.3 below. Overall, these points demonstrate the high potential value of an efficient knowledge sharing mechanism.

4.2.2. Knowledge representation

There were some criticisms of the generic nature of (some of) the information provided (P39). Related to this, there were suggestions for improvement, such as more detailed, precise, and updated information (P27, P37, P40) and more specific instructions (P40). This suggests that knowledge representation should be comprehensive and precise, focussing on providing detailed instructions or information where necessary. Indeed, participants were concerned that some knowledge would be too complex for the system; for example,

'Problems at the production line can be more complex and have complex solutions'; 'Mechanical problems can be hard to describe. pictures might be better'; and 'It will be challenging to match similar problems using text only as people can describe them in many different ways' (P8*). Technically, this was a significant challenge when using keyword-based search and intent-based assistants; however, semantic search and LLM-powered RAG can handle divergent phrasing much better. Other operators believed a simple approach would be sufficient; for example: 'Error code and description would be enough information to save' (P24), and P22 reiterates that 'Error codes are the single most important factor to match to existing issue solutions'. However, P25 points out that 'Sometimes there are random issues that don't have specific fault codes', demonstrating a need for alternatives. These wishes reflect a desire to minimise interaction time with the system. Therefore, a tension exists between the perceived effort required to document knowledge and the value of that knowledge for sharing purposes.

4.2.3. Knowledge maintenance

The comments discuss the necessity and challenges of knowledge maintenance from several perspectives. P40 states that 'If the system is not kept updated, it becomes useless' and P24, whilst reflecting on previous tools for KM, mentioned that 'The biggest challenge is keeping it updated. It takes so many man-hours, and you just don't have those'. Additionally, participants raised concerns about not having sufficient data, such as P8*, who stated that '[The system] might not give good recommendations if it's only half full of knowledge' and '[The system] needs more data'. There are also concerns about how to share knowledge clearly if there are many (potentially conflicting) entries in the knowledge base to share (P8*).

Some resistance to participating in the upkeep of knowledge is evident as operators reported that uploading documents is not their responsibility (P40). There were also concerns about how the database will be connected and updated with instructions and documents (P8*), and the need for a proper process for approving knowledge updates to maintain high quality, as stated here by P8*: 'We will need to have a proper process for the settings change confirmed by an engineer'. Comments on the existing KM situation at the factory portray the challenges in keeping the knowledge updated, as discussed in Section 4.3.4.

4.3. Adoption and change management

This section addresses the socio-technical challenges related to the adoption process of cognitive assistants

for knowledge sharing. This includes a detailed account of the differences between the operator and management perspective, a key component of the research question.

4.3.1. Operator support and assistance

The comments suggest that operator support and assistance are crucial for adopting the CA. P29 highlights the importance of awareness and training for operators, which could potentially hinder adoption if not adequately provided. P37 also emphasises the need for operators to know how to use the system correctly to improve adoption.

4.3.2. Operator acceptance and adoption challenges

The comments reveal several challenges to operator acceptance and adoption of the CA. Managers are concerned that 'If two or three bad experiences occur, then [operators] might be discouraged from using it'. Some participants refer to the need for full adoption; for example, P30 suggests that adoption could be hindered if all employees don't accept the system, indicating a need for broad buy-in. Furthermore, P38 implies that older operators' usage could be a potential barrier to adoption, and P34 also points out that '[User acceptance] Depends on digital literacy of operators', suggesting that age and technological proficiency may play a role in adoption. However, P34 clarifies that they have no concerns about adoption. P35 reveals privacy concerns as a potential obstacle to adoption, with one operator refusing to participate due to discomfort with human tracking, as discussed in Section 4.4.1. This remains a critical challenge for adoption, especially if rogue operators sabotage the system. Other potential barriers to adoption include poor performance (4.5.2) due to lack of details and out-of-date knowledge (4.2.2), poor accessibility (4.5.1), lack of cultural importance for knowledge sharing (4.3.4), and a high perceived effort to record knowledge (4.2.2).

A few participants referred to failed attempts to introduce new documentation tools; for example, P23 points stated: 'We bought a bunch of tablets 7 years ago with the intention of using them for documentation, but they were never adopted' and that current documentation tools are poorly used (see Section 4.3.4 for more details.) Although we do not know the reasons behind the failed adoption of tablets for documentation, it does reflect how challenging introducing a new KM system can be, and why both management and operators need to commit.

4.3.3. Management and operator perspectives

Using a CA to share knowledge among operators is largely motivated by the factory management's observations that this was not naturally occurring effectively despite their attempts to stimulate it. This is reflected in the following statement by P3, an operator: 'Management has accused me of not sharing my knowledge with others. However, it was due to a miscommunication about what the problem was about'. The comments generally indicate a range of perspectives between management and operators regarding knowledge sharing practices. Some operators, especially those who were very distrustful of management, strongly believe that knowledge is already sufficiently shared and do not appreciate the suggestion that they are hiding it. Regarding documentation, some operators stated that operating the production line is too intensive, leaving insufficient time for reporting. From the management's perspective, they observe operators who are frequently idle and not proactive in sharing their expertise. Regardless of the truth, we observed that the careful introduction of a knowledge sharing tool such as a CA is key to reconciling the differing perspectives. For example, by framing the CA as a support tool for operators to avoid the impression that experienced operators are being forced into sharing their knowledge.

Indeed, many of the managers believed that the CA would need to be carefully introduced to avoid rejection from the operators. P29 suggests the need to sensitise operators and explain the system's advantages, implying a potential gap in understanding or perception between management and operators. P38 and P39, both operators, suggest that perspectives may vary depending on the individual, with some being indifferent to the differences between management and operators. P1 reveals a desire for operators to be heard by upper management, whereas others are satisfied with the influence operators already have and feel insulated from the desires of (upper) management through, for example, the workers' council (P32). That being said, several operators mentioned that the management implements new tools without fully committing. This has led to distrust toward management regarding whether they would commit to the CA. Conversely, the management believes that the tools they introduced work sufficiently well, and the operator's laziness is the problem. Overall, our experience is that most operators were open to new knowledge sharing tools introduced by management as long as they were involved in the process, the benefit was clear, and the management invested sufficiently in delivering a reliable and useful tool.

Regarding the knowledge sharing interactions with the CA, we observed differences in requirements

between the operators and management. Management was mostly concerned with collecting sufficiently detailed knowledge and ensuring operators used the tool to encourage standardised best practices. Conversely, the operators were concerned with ensuring the interactions were as short as possible, minimising the required details and only using the tool when absolutely necessary. While (upper) management tended to focus on an idealistic view of capturing every detail of an operator's knowledge about a problem, the operators tended to seek out the minimally required detail for describing and retrieving problem solutions. Concurrently, the vision of (upper) management was that all operators follow a single 'best practice', whereas some operators preferred to follow their own intuition and personal strategies, at least when it came to tuning the parameters of the production line. When it came to solving critical problems that stopped production, operators were more open to learning from others. Overall, we recognised a stark difference in perspectives between management and operators, both with valid concerns. These factors demonstrate the complex socio-technical challenges of introducing a tool for sharing knowledge.

4.3.4. Integration with existing systems and practices

The comments suggest that integrating the CA with existing systems and practices is a key consideration. P1 revealed that some operators use WhatsApp for internal communication, suggesting that operators were not satisfied with official communication channels. That being said, the operators still used the company-issued phones to request help from technical services or expert operators, many of whom were more senior. The official policy for resolving problems at the production line consisted of several steps of escalation, beginning with the present operators trying to solve the problem themselves. If unable to solve the problem, operators can request help from other (expert) line operators, shift leads or technical services. If a CA were introduced, it could be positioned as the first escalation step, allowing (novice) operators to elicit support before needing to call on their fellow operators, who are likely busy with other tasks. Most operators supported the idea of sharing their knowledge to support each other in this way, especially less experienced operators and those involved with training operators. However, it remains to be seen if the operators would be willing to maintain the effort for a prolonged period of time.

At the production line, we noticed a pattern where operators frequently skipped or improperly executed tasks they deemed non-essential. Operators revealed

that some new procedures are not always followed, such as autonomous maintenance procedures (small maintenance tasks that operators are expected to routinely conduct): '[The autonomous maintenance] cards are frequently turned around without the task being completed' (P3). Furthermore, there is evidence that operators do not properly input information into existing documentation systems, as stated by P2: 'Because the problems are so poorly documented in [issue reporting system], we don't have a good overview of the problems' (P2) and 'Existing digital report system is frequently unused' (P5). When we asked the operators why they were not using the issue reporting system, we frequently heard that it was unreliable, suggesting that the operators did not see the point of putting in effort to use something that could not be relied upon. Furthermore, the upkeep of standard working procedures has not been performed as mentioned by P24 here: 'No, no, but in fact, the current standard working instructions are from 2003 or 4 when we started with it. It has become such a tangled mess. We really want a different system for it. It's all still in Excel'. All of this points to a lack of investment in knowledge management systems from the management and a lack of engagement in responsibilities that are perceived as non-essential to operators, possibly reflecting a work culture where operators are disinclined to conduct tasks beyond the operation of the production line. Interestingly, P33 mentioned that 'Previous non-digital version existed and was popular but more cumbersome', indicating that engagement with a previous knowledge sharing system used to be higher but was hindered by its unwieldiness. While the introduction of a CA alone will likely not change the underlying culture directly, it could reduce the perceived effort in documentation and improve the perceived benefit of documentation.

4.4. Privacy, safety and ethics

This section addresses the research question by examining perceived risks and ethical concerns related to the use of cognitive assistants.

4.4.1. Concerns over operator privacy and data handling

The comments reveal mixed opinions regarding user privacy and data handling. Some users, as indicated in comments P31, P32, P33, and P34, do not express any privacy concerns, with P33 stating, 'No [privacy concerns as] most of the data isn't personally identifiable'. However, P35 reveals a contrasting perspective, where an operator refused to participate due to privacy concerns and was uncomfortable with human tracking

and covert management oversight. P34 suggests a balanced approach, advising the provision of individual profiles for operators but also cautioning about privacy. This suggests that while some users are indifferent or unconcerned about privacy, others are acutely aware and concerned about how their data is handled and monitored.

4.4.2. Security measures against misinformation or malicious use

There are some concerns regarding the security risks if the system's advice is incorrect, as stated by P27: 'Risks: if the responses are not adequate, you risk safety'. Safety becomes especially critical in inherently dangerous industries such as detergent factories, the context of this study. This raises questions regarding who is accountable for mistakes, especially if the system relies on knowledge shared by other humans, and what level of accuracy is acceptable for the system if human safety is at stake. Furthermore, P33 mentioned a 'Small possibility exists that users may be malicious and enter falsehoods but unlikely'. The above points support the necessity of mechanisms to safeguard the quality of the knowledge, such as knowledge input approval procedures, checking for conflicts with existing knowledge, and identifying outdated knowledge, many of which could also be supported by AI (see also Section 4.6.3). Furthermore, framing the system as a supporting tool, not a decision-maker, could help remind operators that they are still accountable, as suggested by P34: 'Will be useful as a supportive tool but not a solution'.

4.4.3. Operator autonomy and empowerment

Participants emphasise the importance of operator autonomy and see the CA as a support tool, as stated by P33: '[It will benefit operators] as long as it remains a support tool and operators remain the main actor in the work'. P34 suggests giving each operator their own profile to aid knowledge management, indicating a desire for operator autonomy and empowerment. However, the same comment also cautions about privacy. Indeed, as discussed above in Section 4.4.1, P35 mentioned that an operator refused to participate due to privacy concerns about the camera system, underscoring the importance of operator autonomy. On this topic, operator opinions are divided, as P5 states, 'No issues with having cameras and would rather that [management] invest in enough cameras to cover the entire line instead of half-measures', indicating a desire for management to fully commit to a solution that can support operators. Throughout the process, we respected operator autonomy by involving the workers' council



Figure 4. A sticker used to block the stereoscopic camera for anonymous human tracking that overlooks the production line.

in the design process. They approved the work we conducted, and everyone was given the opportunity to ask questions and discuss concerns. However, despite formally approving the deployment of a camera system, at least one operator sabotaged the installment, indicating disagreement indirectly, as mentioned by P21: '[Operators] have placed a sticker over the camera, blocking its view'. (see Figure 4 below). This highlights an ethical dilemma that organisations and researchers may face when introducing AI systems that track humans: operators may not feel safe to oppose the introduction of new tools openly.

4.5. Usability and user experience

This section addresses the research question by presenting perceptions of usability and user experience, which influence both the impact and the challenges of adopting cognitive assistants.

4.5.1. User interface design and accessibility

Operators highlighted the importance of a user-friendly, accessible system that caters to diverse user needs. Operators expressed a desire for a more intuitive information display, with one suggesting to 'Display the attachments after the main text for easier reference' (P28). Similarly, several suggestions highlight the shortcomings in the systems NLU, for example, requests for recognising keywords to provide relevant suggestions and streamline communication (P8*). The use of visual aids was also highlighted, with suggestions to 'use more pictograms and pictures to speed up understanding' (P8*). Language localisation was another aspect that users found important, with requests for system versions in other (local) languages. Mobile accessibility

was appreciated, but concerns were raised about the use of personal devices, with one user questioning, 'Are we not going to force people to use their phone for that?' (P24). These comments highlight the need for a system that is easy to navigate, visually intuitive, available in local languages, and providing company devices.

4.5.2. System performance and intelligence

The topic of reliable and efficient user interaction is frequently mentioned. Whilst nine ($n = 9$) participants are positive regarding the system's natural language understanding, eight ($n = 8$) participants expressed concerns or areas for improvement, highlighting the challenges for user satisfaction. Suggestions include the ability to understand user queries based on keywords (P8*), including autocomplete in the chatbot to speed up the performance, and better NLU when asking for assistance during machine troubleshooting (P8*). There were also calls to load more (detailed) data and knowledge into the system (e.g. P27) to improve performance. Despite these concerns, $n = 9$ participants provided positive feedback on the system's performance and usability, suggesting that it is generally well-received and effective in its current form. Two participants appreciated the response speed; for example, P29 stated: 'Quick in responses'. Others highlighted the desire for 'More consistent performance' (P33). The emphasis on reliable, speedy performance is logical, considering one of the main perceived advantages of the system is to speed up problem solving as stated by P37: 'Advantage: speed in problem solving'.

4.5.3. User training

The participant suggests training the operators to use the system effectively; for example, P8* stated that 'Probably a short training or demonstration would be needed to be able to use it properly'. This demonstrates that despite the system supporting natural language interaction and being able to explain how to use it, participants still value formal training to improve user experience. Furthermore, P36 stated that 'Advantages in the future with generational change with adequate training, no particular difficulties' and P8* also mentioned that once operators get used to the system, it becomes easier, suggesting that with some help, the system was easy to use.

4.6. Technical and operational issues

This section details the technical and operational challenges encountered in the deployment and use of cognitive assistants for knowledge sharing.

4.6.1. Network and connectivity issues

The comments indicate significant network and connectivity issues that affected the operations. For instance, P35, and P23 highlight that Wi-Fi and 4G internet within the factory were unreliable. This disrupted communication and data transfer, crucial for smooth user interactions with the assistant as it relied on an internet connection. Additionally, P21 points out firewall issues that hindered access to certain resources. This demonstrates the importance of a reliable and well-managed Wi-Fi network within organisations considering deploying connected systems such as CAs.

4.6.2. Technical limitations and troubleshooting

We faced several challenges integrating the system with factory systems. P35 suggests that coordinating with the (3rd party) factory IT was challenging due to their frequent unavailability or preoccupation with other tasks. Furthermore, P21 mentions a technical issue where a PC was inaccessible, highlighting the limitation of coordinating with multiple technical parties to integrate systems.

4.6.3. Dependence on accurate and complete data input

The comments highlight challenges related to the 'cold start problem': providing benefits and driving engagement if the system's knowledge base is empty. Indeed, during evaluations, P8* indicated that the system needs more data, confirming this issue. Motivating operators to share their knowledge is a considerable challenge, especially initially. Furthermore, once the system has reached a point where it provides perceptible benefits, it will still need to be maintained and kept up-to-date. Factory management also required that the CA include a validation step so that an expert operator can approve all incoming knowledge. Furthermore, P24 emphasises that the biggest challenge they faced with previous KM systems was keeping them updated due to the significant man-hours required. This suggests a need for more efficient data input methods or automation to ensure the system has accurate and complete data. Additionally, a mechanism is required to remove and update existing knowledge as the machines and products can change over time; P25 states, 'Yes, sometimes we have products that change their speed or the machine is modified, allowing faster operation'. This underscores the importance of having an efficient mechanism for keeping a knowledge base up-to-date and approving incoming knowledge.

In sum, operators primarily perceive CAs as a way to speed up troubleshooting and reduce dependence on scarce experts, while managers emphasise

standardisation and performance gains. A thematic analysis of the comments from reveal three main socio-technical risks (privacy, safety of advice, and cultural resistance) and four practical challenges (knowledge upkeep, authoring burden, infrastructure, and change management).

5. Discussion

The domain of computer-supported knowledge sharing has evolved significantly over the decades, aligning with technological advancements, a focus on labour relations, and shifts in work practices. Originally, computer-supported knowledge sharing focussed on centralising expert knowledge through the 'repository model' in the late 1980s. This model primarily relied on forming extensive databases that document expert insights. However, this strategy struggled with situational variability as knowledge is often highly contextual and not easily generalised across different scenarios. Furthermore, the requirement for experts to codify their knowledge into repositories imposed a significant authoring burden, detracting from their primary tasks and responsibilities.

Subsequently, the field shifted away from static repositories towards facilitating direct exchanges between individuals, aiming to capture the nuanced and dynamic nature of knowledge. However, scalability remained a challenge, and the effectiveness of these interactions hinged on the availability of both parties, which is not always feasible in a high-paced shift-based production environment.

Our research aligns with the paradigmatic shift towards cyber-physical systems as intermediaries for knowledge sharing (Hoffmann et al. 2019). These systems, particularly through advancements in AI and live data processing, aim to manage the complexity of real-time knowledge sharing. This leads to critical discussions in various domains on how to share knowledge from a human to machine and visa versa (Abubakar et al. 2019; Lista Rossetti et al. 2023; Schniederjans, Curado, and Khalajhedayati 2020). Although, still underrepresented, our work contributes to the small growing body of HCI and CSCW work in manufacturing such as (Cheon, Schneiders, and Skov 2022; Gerbracht et al. 2024; Hoerner, Schamberger, and Bodendorf 2022; Hoffmann et al. 2022; Mörike 2022; Wurhofer et al. 2018). Our work goes beyond prior work on knowledge sharing in manufacturing such as (Hannola et al. 2020; Hoerner, Schamberger, and Bodendorf 2022; Hoffmann et al. 2022) by learning from evaluating fully functioning CAs in real-world production environments over a period of two years and integrating recent technology advancements, namely LLMs.

Our findings suggest that while (LLM-powered) CAs effectively bridge gaps left by earlier paradigms, such as scalability and minimising the authoring burden, they also introduce new dimensions to consider, such as the potential perceptions of surveillance, and the reliability of advice delivered by AI. Through this discussion, we situate our work within the broader HCI and CSCW computer-supported knowledge sharing research, highlighting the challenges and lessons learned in the form of design guidelines (G#).

5.1. Cognitive assistants excel at sharing the solutions to emerging issues

In this work, the CAs were perceived to be effective at sharing knowledge, especially when acquiring knowledge from experienced operators to support novices. Both management and operators recognised that emerging issues at the production line warranted rapid dissemination of solutions that currently did not happen effectively. As production lines become increasingly complex, digitised, and adaptable, operators spend a considerable amount of time and cognitive resources on error handling, in line with findings of prior work (Wurhofer et al. 2018). Concretely, a few hours of downtime could be saved on the production line weekly if solutions to emerging issues are shared effectively. To maximise this benefit, it is crucial to ensure that information can be found quickly and the knowledge base can be easily updated or modified, leading to the following guidelines: Design CAs to facilitate rapid information retrieval to support timely resolution of issues (G1) and Design CAs to enable quick modification and additions to the knowledge base (G2). Overall, participants universally agreed on the benefits of sharing solutions to issues such as machine malfunctions or problems caused by human error to resolve the issues quicker. In reference to the research question, this benefit forms the primary positive impact that introducing a CA for knowledge sharing can have.

In contrast to sharing solutions for recurring problems, the topic of sharing knowledge about tuning the machines to achieve higher performance is controversial for some operators. This is especially true for those who are comfortable with operating at lower, stable speeds. These operators were skeptical about the benefits of knowledge sharing on machine fine-tuning as they apply unique strategies that others cannot easily adopt and are highly situated. Machine tuning is a complex and iterative process that is affected by many factors, such as ambient temperature and fluctuating raw material properties, making it more complex to share, aligning with observations by prior HCI and CSCW

research (Hoffmann et al. 2022). On top of this, according to management, some operators attempted to hide their knowledge to maintain their position of power. As such, we observed resistance to sharing valuable knowledge on machine tuning for various reasons.

While some experienced operators see no benefit in learning from others concerning machine tuning, this contradicts the management perspective, who observed significant differences in the production speeds attained by different experienced operators and were trying to encourage them to learn from each other. However, some operators still recognised the benefit of sharing some information about machine tuning, such as attainable target production speeds, the accompanying parameters, and some expert tips. Ultimately, in this action research study, we shifted focus to supporting the operators during issue resolution instead of machine tuning, as the benefits were clear to all. Thus, we would advise to collaborate with operators to identify the scenarios where sharing knowledge via a CA can result in significant and perceptible support for the operators (G3).

To overcome adoption issues due to a ‘cold start’ whereby the knowledge base is sparse or empty to begin with (Chen et al. 2023), we prepopulated the knowledge base of the CAs with knowledge we extracted from existing issue reports. We also used the issues captured from the first phase of the study for the later phases. We observed that while this initial population of the knowledge base was only a small subset, it provided sufficient coverage for operators to get a feel for the potential of the CAs. This was aided by the CAs abilities to infer possible recommendations based on similar issues even if a perfect match was not found. Thus, we suggest pre-populating the CA’s knowledge base with knowledge elicited from existing issue reports or interviews (G4) and to use techniques to infer recommendations over sparse knowledge when exact matches aren’t found, such as reasoning over knowledge graphs (Du et al. 2022; Wang et al. 2019) or using LLMs to infer suggestions based on a selection of similar problems.

5.2. Considering the authoring burden of knowledge and the relevance for other operators

While CAs may reduce the authoring burden of knowledge sharing somewhat; it still remains a core socio-technical challenge. Several participants brought up the challenge of motivating engagement from operators. While improving accessibility and keeping dialogues short can minimise the perceived effort of using the CA, the perceived benefit of sharing knowledge to help others remains a central barrier to adoption

aligning with prior work (Hannola et al. 2020; Hoerner, Schamberger, and Bodendorf 2022). Looking at the factors that affect technology acceptance (Davis 1989), perceived benefit and perceived effort as opposing forces that will impact the attitude toward using the technology. External factors can also influence adoption, such as incentives, cultural norms, and voluntariness to help colleagues. In this study, we observed the strongest motivation to share knowledge among the ‘expert’ operators who were already involved in training novice operators as part of their job.

Going deeper into the tensions surrounding the balance of providing and receiving knowledge, we assume that more effort invested in eliciting and capturing knowledge will result in a more comprehensive knowledge base that can be used to help others. Understandably, the balance may shift depending on many factors, such as its potential impact, how time-critical it is, if it’s something only one expert has discovered, how broadly applicable it is, how situational it is, how readily it can be shared, etc. For example, a solution to a recurring and high-impact problem that only one operator has discovered should be immediately captured and shared, even if it takes considerable effort. Additionally, the act of sharing knowledge is in tension with conducting the primary task of operating the production line, echoing previous work by Hoffmann et al. (2022). It is also important to consider the ephemeral nature of knowledge as the complex system changes and new knowledge is discovered. As such, organisations could consider a prioritisation framework to guide operators on how much effort should be invested in sharing their knowledge on a case-by-case basis. This prioritisation can be automated, using frequently asked questions received by the CA to determine relevance and production urgency. Therefore, to mitigate the burden of knowledge sharing on operators, the value of that knowledge must guide the prioritisation of what to share and when (G5)

Maintaining the temporal relevance and quality of the knowledge base is a key concern from both management and operator perspectives. The managers involved stipulated that all new knowledge entries should be approved by designated ‘expert’ operators, similar to mechanisms deployed by Hoerner, Schamberger, and Bodendorf (2022), providing some level of quality control. In this study, management was enthusiastic about using LLMs to validate the knowledge inputs by checking logic and consistency with existing knowledge. The validation output can be used by the operator to improve the report before submitting it for human approval.

Although not implemented in our work due to the anonymity policies we agreed with operators, other

work has made the author of entries visible to others (Hannola et al. 2020), which could support the validation and knowledge application process but also introduce potential biases concerning the quality of specific operator's knowledge (Yang, Yuan, and Wang 2019). Additionally, the validity of entries should be continuously evaluated, for example, by enabling operator feedback or by identifying conflicts between entries and safeguarding against outdated knowledge (G6). Overall, several strategies can be employed to maintain a high-quality knowledge base while considering the potential for bias, the burden on knowledge maintenance, and resulting improvements.

5.3. Educating the local stakeholders is key

One of the social challenges faced is related to the adoption of new technologies, such as CAs. Operators who fundamentally disagree with the benefit of sharing some knowledge may avoid using the CA. Indeed, assuming that the CA can support operators in doing their job better, the operators that are open to using the CA and can effectively interact with it will be at an advantage. For new operators, this means they can become more independent of their human mentors more quickly, and more experienced operators can benefit from the collective knowledge of their colleagues. In contrast, the perceived value of the operators not using the CA may be reduced, for example, because they reject it fundamentally or have difficulties interacting with it. This highlights the importance of considering accessibility and, where possible, facilitating full operator adoption, for example, by involving them throughout the development process and solving problems that are important to them (G7).

During our study, many operators demonstrated a deep pride in the knowledge and expertise they apply in their work. It becomes critical, therefore, to frame technology, such as a CA, as a tool that supports rather than replaces their skill and judgment. This notion is in harmony with Shneiderman (2022) who advocates for 'humans in the group; computers in the loop', stressing that CAs should enhance rather than substitute human decision-making. Emphasizing that operators are still responsible for their decisions even with CA support could help them stay critical and informed about the advice they receive from the system (G8). Additionally, communicating the strengths and weaknesses of the CA, such as being aware of potential 'hallucinations' in responses from Large Language Models (LLMs) and the influence of leading questions, is crucial for productive interactions (G9).

Yet, similar to prior research, we observed that workers' councils and operators often do not know

enough about the technology to engage deeply in discussions concerning new AI systems (Gerbracht et al. 2024). On top of this, upon reflecting on our own experience with the workers' councils, we recognise that they were only involved upfront during the project initiation when many details were still fuzzy. In contrast, after the project was approved, they were no longer involved. Considering the innovative and explorative nature of the project, it could have benefited from more continued input from the workers' council to safeguard the operator's interests and ensure they remained well-informed, as demonstrated by Crabtree, Tolmie, and Knight (2017) and R. Zhang et al. (2021). We also note that despite our efforts to explain the CA architecture and capabilities in our interactions with operators and posters we put up at the production, we observed that many operators only fully grasped the concept and its benefits once the full system was deployed. This could be attributed to the complexity and novelty of the system which warrants continuous efforts to educate relevant factory stakeholders about the system and its impacts, such that they can make informed decisions about its adoption (G10).

5.4. Co-evolving local infrastructure and mitigation plans for technical issues

In addition to the social challenges, deploying CAs at the production line faces many practical and technical challenges that affect the system and our research. The factories in this study used numerous separate digital systems such as for planning, maintenance activities, quality control, productivity, and machine documentation. Many of these systems were considerably outdated and incompatible with each other. As a result, considerable effort was needed to connect them to the CA, sometimes requiring the involvement of third-party suppliers. For example, some of the underlying infrastructures, such as databases with API access, needed to be deployed before the research could proceed. As such, considering the effort involved, relevant stakeholders should discuss whether connecting specific systems to the CA is worth it, considering factors such as added value, reliability, the investment needed, privacy implications, and upcoming system changes (G11). As W. K. Edwards, Newman, and Poole (2010) and Martinez-Maldonado et al. (2023) have highlighted, developing HCI infrastructures that adapt alongside work practices is not just beneficial but essential for the sustainable implementation of technological advancements in any setting.

The poor availability and quality of the data were compounded by poor (wifi and mobile) network

stability in the factory, meaning the availability of live data could not be guaranteed. As a result, operators sometimes needed to wait 30 seconds for a response or could not receive the support they requested. Therefore, we advise that companies should prepare reliable and robust data infrastructures before deploying CAs (G12), for example, by providing reliable wifi and investing in consolidating (live) data in centralised databases on modern servers. Acknowledging that high-quality, live data will not always be available, designers and developers of CAs should design cognitive assistants to be robust to missing or low-quality data (G13). For example, mitigation strategies could include proactively asking the operator to provide critical information describing a problem (e.g. pressure levels) if the CA needs it to give reliable advice.

Accessible AI for Support, not Surveillance A key social challenge is related to tension between operators and management regarding the reporting of what is happening at the production line. While existing methods of reporting issues were often poorly executed, operators were not very motivated to improve them. Several operators mentioned that the current system used for tracking productivity were not very reliable as they did not accurately measure production speed. Similarly, operators complained about the reliability of the issue reporting system. Some operators used this to justify why they did not report all issues or pay much heed to the reported production statistics. Thus, when a manager tries to investigate why production was low for a shift, the operators claim that the data is inaccurate or does not represent the actual situation. This situation is similar to operator behaviour observed by Mörike (2022), where operators deliberately entered incomplete data into a digital terminal so that management did not have a complete picture of their work. Consider that operators may resist systems that provide managers with more accurate and timely information on production processes (G14). Thus, understanding the origins of operator resistance to system changes is vital for designing CSCW and HCI solutions that will be adopted and used effectively.

Although most operators expressed concerns about privacy, particularly regarding surveillance by managers and being unfairly evaluated on work performance, the majority agreed to allow sensitive data collection, such as their location. The main requirement was that access was restricted, particularly from management, highlighting the necessity that the benefits of data collection should be clearly beneficial to operators and not just for surveillance, echoing findings of Das Swain et al. (2024) that found information workers perceived tracking as pure surveillance when they were not presented any

consumable insights. Surprisingly, the operators we interviewed went as far as to express favouring a more intrusive tracking system that could enhance the system's effectiveness over half-measures that did not benefit them. Thus, any implementation of tracking technologies must carefully balance utility and privacy while including operators in the decision process (G15).

A major challenge when involving operators in privacy discussions is catering to widely different perspectives. In our study, the CA is context-aware via an operator tracking system. The tracking system consisted of multiple stereoscopic cameras and a computer vision model to create an anonymous 18-point skeleton of operators along the production line. The tracking data was not personally identifiable, and factory employees could not access it. The workers' council and all the operators we spoke to approved the tracking system. Furthermore, all operators were informed about the tracking system via announcements and posters along the production line. Despite these efforts, the tracking system was sabotaged by an operator. Despite pressure from factory management to fix the tracking system and continue with the research, we interpreted the operator's actions as an implicit rejection of the tracking system and disabled it. Thus, beyond respecting the legal side of privacy, organisations will need to consider when to force the implementation of human tracking for CAs, especially if the majority of the workforce would prefer the additional utility that context awareness could provide, such as automatically capturing knowledge (Yang, Yuan, and Wang 2019) or suggesting responses.

Beyond context-aware suggested responses, several other factors can impact the accessibility of interactions with the CA at the production line. These factors include familiarity with CAs, difficulty reading text on a smartphone screen, physical access to the smartphone with the CA, noise that may impact the accuracy of speech-to-text and audibility of text-to-speech, (local) accents or jargon affecting speech-to-text and NLP, and existing knowledge on the abilities and limitations of AI technology, such as LLMs. Indeed, aligning the design of the CA with the needs of operators is a crucial area of research (Fox et al. 2020). The challenges of handling noise, accents, and local jargon when using speech-to-text remains an open point but beyond the scope of our work. In this study, we observed that providing (context-aware) suggested responses for the operators to click on instead of type responses was effective at improving accessibility. Furthermore, using advanced NLP technologies, such as LLMs, made the CA easier to use as it became more capable in understanding the operator's queries. The final version of

the CA we deployed was also accessible in computer browsers instead of only through a smartphone app, improving accessibility for some of the older operators who had trouble using the smartphone. These integrations underline the importance of designing technology that is adapted to real-world factory contexts and the capabilities of the operators (G16).

5.5. Implications for researching and deploying cognitive assistants in factories

Navigating the tensions between operators and management regarding AI-mediated knowledge sharing is key to the success of CAs in factories. While introducing a CA may facilitate the sharing of knowledge, the organisation's culture regarding collaboration and trust will significantly affect its positive impact. This was evident in the operators' preference for autonomy and personal strategies, whereas management was more focussed on standardising practices and boosting efficiency. Our study exposed several underlying trust and cultural issues between operators and managers that appeared to exacerbate the misalignment of these priorities.

As these issues are not unique to using CAs for knowledge sharing, we can look to prior literature on building organisations that encourage knowledge sharing. For example, emphasising 'knowledging' over transactional, static knowledge transfer (Cheuk and Dervin 2011), providing sufficient justification from managers when introducing new processes (Olander et al. 2016), encouraging a mastery culture over a comparison culture (Lee et al. 2022), and portraying the CA as a way to share learnings autonomously (Monks et al. 2016), all support knowledge sharing based on mutual trust and a common purpose rather than duty (Oleksa-Marewska 2020). Managers could introduce low-key integration of knowledge sharing into daily briefings, for example, by mentioning specific knowledge sharing needs or giving positive feedback on prior knowledge sharing (Ellström and Ellström 2014; Israilidis, Siachou, and Kelly 2020; Sijbom et al. 2025). While face-to-face interactions and informal conversations are traditional strategies to build trust and make knowledge sharing more enjoyable (Oliveira et al. 2023; van Greunen, Venter, and Sharp 2021), future research could explore if similarly positive experiences can be (partially) achieved when sharing via a CA. Thus, we can conclude that organizations considering introducing CAs for knowledge sharing should also create a work culture that is conducive to knowledge sharing (G17).

As the management initiates the decision to bring in new technology like CAs, the operators can perceive this

as a means of control and surveillance. As discussed by prior studies, the early and continued involvement of operators and workers' councils can facilitate the adoption of new technology in factories but did not go into the specifics on how (Cheon, Schneiders, and Skov 2022; Meneweger et al. 2015). The technology investigated in this study, (LLM-powered) CAs that support knowledge sharing, pose some unique challenges that warrant extra attention, such as the difficulty imagining the impact on work, the authoring burden on (expert) operators, focussing on knowledge that operators understand the benefits of dissemination, understanding the limitations of LLMs and the possibility of feeling surveilled by management. Unlike prior work that required managers to mediate all interactions with operators (Cheon, Schneiders, and Skov 2022), we were able to involve them separately in most cases. Thus, we could maximise the openness and directness of the responses we received by positioning ourselves as external researchers (i.e. not 'agents' of the management), building a relationship with participants over time, talking to operators without management present whenever possible, and maintaining the anonymity of participants (G18).

From our work, we recognise that the most crucial topic to involve operators is identifying what knowledge would be beneficial to share and how it can be expressed and retrieved (G19). Tackling this together requires clear communication and education about what the technology can and cannot do. In this study, we witnessed the importance of respecting the operator's expertise and autonomy, making sure that operators feel their contributions through the CA are valued, not just as data for the system but as important knowledge that improves the workplace. For example, by focussing on sharing knowledge about resolving concrete machine issues as opposed to the less tangible machine tuning knowledge, more operators perceived the CA to be beneficial. This is especially important as many operators were apprehensive of introducing (yet) another tool and the potential consequences thereof, similar to observations made by Meneweger et al. (2015). By ensuring everyone understands and sees the benefits of CAs, factories may reduce resistance and may increase both productivity and job satisfaction (G20).

Regarding social interactions, while AI allows for quick access to information, it could also lead to fewer face-to-face interactions, affecting the social environment. Additionally, increasing dependence on AI means operators must develop new skills, particularly in using advanced digital tools, instead of relying solely on experiential knowledge. This shift could also affect

operators' well-being, balancing between the satisfaction from autonomously or collaboratively solving complex problems and the ease of AI help as previously discussed by Wurhofer et al. (2018). We observed that the operators were proud of their expertise, and introducing the CA to capture and share this might affect the satisfaction they get from work.

While using a CA for knowledge sharing promises to enhance operational efficiency, it is vital to balance using the AI tool for all information needs versus the irreplaceable value of human knowledge sharing, if available (G21). In fact, several operators preferred learning from human operators, similar to findings in prior work (Hoffmann et al. 2022). Yet, a key benefit of the tool is that it could help with problems when a knowledgeable human colleague is unavailable or unable to help.

The long-term effects of using AI for sharing knowledge in factories are likely complex, covering operational, psychological, and social aspects. Using AI for this purpose might improve operational efficiency and make the knowledge management system within organisations more dynamic and responsive. This may change how organisations are structured, making operators more independent from other (expert) operators. However, there may be fewer opportunities to share more subtle, tacit knowledge that is more challenging to share through technology, emphasising the need for research in these areas.

A well-established concern in HCI/CSCW literature is that AI decision support may lead to erosion of human critical thinking Bućinca, Malaya, and Gajos (2021), Vasconcelos et al. (2023), Y. Zhang, Vera Liao, and Bellamy (2020) and Harbarth et al. (2024). While the participants in our study did not voice this concern explicitly, they did emphasise the importance of framing the CAs as a support tool. We believe three contextual factors explain this finding: (1) the CAs were never fully adopted on the production line; (2) the operator culture of being proud of their expertise and abilities to solve problems; and (3) the participant's concern on the risks of inaccurate AI output while led to the study's deliberate framing of the CA as a support tool, reinforcing operator accountability. Indeed, this concern has become more relevant in the latter half of the study, with the introduction of LLMs and thus, the occurrence of very plausible but hallucinated advice. While skeptical expert operator might be able to spot LLM hallucinations, novice operators may not have sufficient experience to do so. As such, additional efforts should be made to avoid undesirable outcomes when using LLMs. For example, for this study, we clearly communicated the limitations of the CAs, and

incorporate mechanisms to double-check recommendations. These long-term concerns highlight the importance of longitudinal studies at factories to evaluate the impacts and mitigation strategies. This future research will need to navigate challenges related to the safety and productivity concerns in factories when piloting new technologies (Cheon, Schneiders, and Skov 2022).

5.6. Limitations and future work

In conducting this study, we faced numerous challenges tied to fieldwork in operational settings where financial, health, and safety concerns are paramount. Consequently, the tested prototypes were never utilised for extended periods during actual work activities; instead, we relied on a series of brief user tests, each lasting about 30 minutes and in the case of phase 1, up to 8 hours. While this approach facilitated immediate feedback and iterative improvements, a more prolonged deployment would likely uncover deeper insights into how users gradually adapt to and integrate CAs in their daily workflows. Despite the limitations, the high technology readiness levels (TRLs 7–8) of the prototypes, testing within the real-world factory environment, and our many contact points with participants over two years provide strong support for the ecological validity of our findings.

Furthermore, we acknowledge that case studies inherently produce context-specific insights, particularly in complex socio-technical environments such as factories. The findings and implications drawn from our work are closely tied to the unique organisational culture, workflows, and infrastructure of the studied factories. As such, caution should be exercised when generalising these results to other settings, and future research should investigate how similar interventions perform across diverse manufacturing contexts.

This study underlines the importance of socio-technical alignment when implementing AI-driven systems like CAs. Over the course of this longitudinal study, we shifted toward sharing knowledge about problem-solving as all operators recognised it as important, and we identify opportunities to support the process with CAs. The operators were less enthusiastic about sharing knowledge about other areas, such as machine setup and tuning, due to the situatedness, temporal nature, and personal aspects of the knowledge. However, this knowledge is valuable to factories and warrants focussed research to explore the challenge of sharing this more complex knowledge. Furthermore, HCI and CSCW research should explore the long-term impacts of deploying such technologies, focussing on how they reshape organisational structures, job roles, and the balance between digital- and human-driven problem-

solving processes. This includes addressing the challenges of maintaining critical thinking and decision-making skills among workers in increasingly AI-supported environments and exploring strategies to balance technological assistance with essential human capabilities.

6. Conclusion

This study illustrates both the potential and the complexity of introducing Cognitive Assistants (CAs) for knowledge sharing in factory settings. CAs can speed up problem-solving and help new operators become more independent, but their success depends on more than just technical performance. Key risks include privacy concerns, the reliability of advice, and cultural resistance to changing established work practices. Practical challenges such as keeping the knowledge base current, minimising the effort required from operators, and ensuring robust technical infrastructure also play a major role.

The existing culture of collaboration and knowledge sharing within a factory will greatly influence how successful CAs can be. Addressing these issues requires involving all stakeholders early and throughout the process. Transparent communication about how CAs work and how data is handled is essential for building trust. It is also important to recognise that not all knowledge is equally easy to share. Operators are generally open to sharing solutions to recurring problems, but may be less willing to share personal strategies or context-specific tips.

Overall, successful adoption of CAs in factories depends on finding the right balance between technological efficiency and respect for human expertise, autonomy, and workplace culture. By focussing on both technical and social factors, organisations can better realise the benefits of CAs while addressing the concerns of those who use them.

Notes

1. <https://www.anthropic.com/claude> – last accessed October 11, 2025
2. <https://llis.nasa.gov/> – last accessed October 11, 2025

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work was supported by the European Union's Horizon 2020 research and innovation program via the project COALA 'COgnitive Assisted agile manufacturing for a

Labour force supported by trustworthy Artificial Intelligence' (Grant agreement 957296).

ORCID

S. Kernan Freire  <http://orcid.org/0000-0001-8684-0585>

T. He  <http://orcid.org/0000-0002-8486-3940>

C. Wang  <http://orcid.org/0000-0001-8213-6582>

E. Niforatos  <http://orcid.org/0000-0002-0484-4214>

A. Bozzon  <http://orcid.org/0000-0002-3300-2913>

References

- Abubakar, Abubakar Mohammed, Hamzah Elrehail, Maher Ahmad Alatailat, and Alev Elçi. 2019. "Knowledge Management, Decision-Making Style and Organizational Performance." *Journal of Innovation & Knowledge* 4 (2): 104–114. <https://doi.org/10.1016/j.jik.2017.07.003>.
- Ackerman, Mark S., Juri Dachtera, Volkmar Pipek, and Volker Wulf. 2013. "Sharing Knowledge and Expertise: The CSCW View of Knowledge Management." *Computer Supported Cooperative Work (CSCW)* 22 (4–6): 531–573. <https://doi.org/10.1007/s10606-013-9192-8>.
- Alavi, Maryam, and Dorothy E. Leidner. 2001. "Review: Knowledge Management and Knowledge Management Systems: Conceptual Foundations and Research Issues." *MIS Quarterly* 25 (1): 107–136. <https://doi.org/10.2307/3250961>.
- Angulo, Cecilio, Alejandro Chacón, and Pere Ponsa. 2023. "Towards a Cognitive Assistant Supporting Human Operators in the Artificial Intelligence of Things." *Internet of Things* 21:100673. <https://doi.org/10.1016/j.iot.2022.100673>.
- Balayn, Agathe, Gaole He, Andrea Hu, Jie Yang, and Ujwal Gadiraju. 2022. "Ready Player One! Eliciting Diverse Knowledge Using A Configurable Game." In *Proceedings of the ACM Web Conference 2022 (WWW '22)*, 1709–1719. New York, NY, USA: Association for Computing Machinery. <https://doi.org/10.1145/3485447.3512241>.
- Baldauf, Matthias, Sebastian Müller, Arne Seeliger, Tobias Küng, Andreas Michel, and Werner Züllig. 2021. "Human Interventions in the Smart Factory—A Case Study on Co-Designing Mobile and Wearable Monitoring Systems with Manufacturing Staff." In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems (CHI EA '21)*, 1–6. New York, NY, USA: Association for Computing Machinery. <https://doi.org/10.1145/3411763.3451774>.
- Bang, Yejin, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, et al. 2023. *A Multitask, Multilingual, Multimodal Evaluation of ChatGPT on Reasoning, Hallucination, and Interactivity*. arXiv:2302.04023 [cs.CL].
- Baskerville, Richard L. 1999. "Investigating Information Systems with Action Research." *Communications of the Association for Information Systems* 2 (3): 1–32. <https://doi.org/10.17705/1CAIS.00219>.
- Becerra-Fernandez, Irma, and Rajiv Sabherwal. 2014. *Knowledge Management: Systems and Processes*. New York, USA: Routledge.
- Berger, Patrick, and Joerg von Garrel. 2023. "How to Design a Value-Based Chatbot for the Manufacturing Industry: An

- Empirical Study of an Internal Assistance for Employees.” *Kunstliche Intelligenz* 37 (2–4): 203–211. <https://doi.org/10.1007/s13218-023-00817-6>.
- Bobrow, Daniel G., and Jack Whalen. 2002. “Community Knowledge Sharing in Practice: The Eureka Story.” *Reflections: The SoL Journal* 4 (2): 47–59. <https://doi.org/10.1162/152417302762251336>.
- Boehner, Kirsten, Janet Vertesi, Phoebe Sengers, and Paul Dourish. 2007. “How HCI Interprets the Probes.” In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. New York, USA: ACM. <https://doi.org/10.1145/1240624.1240789>.
- Bradeško, Luka, Michael Witbrock, Janez Starc, Zala Herga, Marko Grobelnik, and Dunja Mladenić. 2017. “Curious Cat–Mobile, Context-Aware Conversational Crowdsourcing Knowledge Acquisition.” *ACM Transactions on Information Systems* 35 (4): 1–46. <https://doi.org/10.1145/3086686>.
- Breque, Maija, Lars de Nul, and Athanasios Petridis. 2021. *Industry 5.0 – Towards a Sustainable, Human-Centric and Resilient European Industry – European Commission*. Technical Report. European Commission.
- Brown, John Seely, and Paul Duguid. 2000. “Balancing Act: How to Capture Knowledge without Killing It.” *Harvard Business Review* 78 (3): 73–80.
- Brown, Tom, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, et al. 2020. “Language Models Are Few-Shot Learners.” In *Advances in Neural Information Processing System*, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, 1877–1901. New York, USA: Curran Associates, Inc.
- Buçinca, Zana, Maja Barbara Malaya, and Krzysztof Z. Gajos. 2021. “To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI in AI-assisted Decision-Making.” *Proc. ACM Hum.-Comput. Interact.* 5 (CSCW1): 1–21. <https://doi.org/10.1145/3449287>.
- Cabitza, Federico, Andrea Campagner, and Carla Simone. 2021. “The Need to Move away from Agential-AI: Empirical Investigations, Useful Concepts and Open Issues.” *International Journal of Human-Computer Studies* 155:102696. <https://doi.org/10.1016/j.ijhcs.2021.102696>.
- Casillo, Mario, 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. New York: IEEE State. <https://doi.org/10.1109/TALE48869.2020.9368339>.
- Castane, G., A. Dolgui, N. Kousi, B. Meyers, S. Thevenin, E. Vyhmeister, and P.-O. Ostberg. 2023. “The ASSISTANT Project: AI for High Level Decisions in Manufacturing.” *International Journal of Production Research* 61 (7): 2288–2306. <https://doi.org/10.1080/00207543.2022.2069525>.
- Chang, Chun-Ming, Meng-Hsiang Hsu, and Chia-Hui Yen. 2012. “Factors Affecting Knowledge Management Success: The Fit Perspective.” *Journal of Knowledge Management* 16 (6): 847–861. <https://doi.org/10.1108/13673271211276155>.
- Chen, Zefeng, Wensheng Gan, Jiayang Wu, Kaixia Hu, and Hong Lin. 2023. “Data Scarcity in Recommendation Systems: A Survey.” *ACM Transactions on Recommender Systems* 3:1–31. <https://doi.org/10.1145/3639063>.
- Cheon, EunJeong, Eike Schneiders, and Mikael B. Skov. 2022. “Working with Bounded Collaboration: A Qualitative Study on How Collaboration Is Co-constructed around Collaborative Robots in Industry.” *Proceedings of the ACM on Human-Computer Interaction* 6:369:1–369:34. <https://doi.org/10.1145/3555094>.
- Cheuk, Bonnie, and Brenda Dervin. 2011. “Leadership 2.0 in Action: A Journey from Knowledge Management to ‘Knowledgeing’.” *Knowledge Management & E-Learning: An International Journal* 3:119–138. <https://doi.org/10.34105/j.kmel.2011.03.012>.
- Colabianchi, Silvia, Andrea Tedeschi, and Francesco Costantino. 2023. “Human-Technology Integration with Industrial Conversational Agents: A Conceptual Architecture and a Taxonomy for Manufacturing.” *Journal of Industrial Information Integration* 35:100510. <https://doi.org/10.1016/j.jii.2023.100510>.
- Cooke, Nancy J. 1994. “Varieties of Knowledge Elicitation Techniques.” *International Journal of Human-Computer Studies* 41 (6): 801–849. <https://doi.org/10.1006/ijhc.1994.1083>.
- Crabtree, Andy, Peter Tolmie, and Will Knight. 2017. “Repacking ‘privacy’ for a Networked World.” *Computer Supported Cooperative Work (CSCW)* 26 (4–6): 453–488. <https://doi.org/10.1007/s10606-017-9276-y>.
- Cullen, Jet, and Alan Bryman. 1988. “The Knowledge Acquisition Bottleneck: Time for Reassessment?” *Expert Systems* 5 (3): 216–225. <https://doi.org/10.1111/exsy.1988.5.issue-3>.
- Das Swain, Vedant, Lan Gao, Abhirup Mondal, Gregory D. Abowd, and Munmun De Choudhury. 2024. “Sensible and Sensitive AI for Worker Wellbeing: Factors That Inform Adoption and Resistance for Information Workers.” In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI ’24)*, 1–30. New York, NY, USA: Association for Computing Machinery. <https://doi.org/10.1145/3613904.3642716>.
- Davis, Fred D. 1989. “Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology.” *MIS Quarterly* 13 (3): 319–340. <https://doi.org/10.2307/249008>.
- Dhuieb, Mohamed Anis, Florent Laroche, and Alain Bernard. 2013. “Digital Factory Assistant: Conceptual Framework and Research Propositions.” In *Product Lifecycle Management for Society (PLM 2013) (Berlin, 2013)*, edited by A. Bernard, L. Rivest, and D. Dutta, 500–509, Vol. 409. Berlin: Springer-Verlag.
- Du, Yuntao, Xinjun Zhu, Lu Chen, Ziquan Fang, and Yunjun Gao. 2022. “MetaKG: Meta-Learning on Knowledge Graph for Cold-Start Recommendation.” *IEEE Transactions on Knowledge and Data Engineering* 35 (10): 9850–9863. <https://doi.org/10.1109/TKDE.2022.3168775>.
- Duhigg, Charles. 2016. “What Google Learned from Its Quest to Build the Perfect Team.” *The New York Times Magazine* 26 (2016): 1–18.
- Dyer, Jeffrey H., and Kentaro Nobeoka. 2000. “Creating and Managing a High-Performance Knowledge-Sharing Network: The Toyota Case.” *Strategic Management Journal* 21 (3): 345–367. [https://doi.org/10.1002/\(ISSN\)1097-0266](https://doi.org/10.1002/(ISSN)1097-0266).
- Edwards, Brett, Michael Zatorsky, and Richi Nayak. 2008. “Clustering and Classification of Maintenance Logs Using Text Data Mining.” In *Volume 87-Data Mining and*

- Analytics* 2008, 193–199. Darlinghurst: Australian Computer Society.
- Edwards, W. Keith, Mark W. Newman, and Erika Shehan Poole. 2010. “The Infrastructure Problem in HCI.” In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Atlanta, Georgia, USA) (CHI '10)*, 423–432. New York, NY, USA: Association for Computing Machinery. <https://doi.org/10.1145/1753326.1753390>.
- Ellström, Eva, and Per-Erik Ellström. 2014. “Learning Outcomes of a Work-Based Training Programme : The Significance of Managerial Support.” *European Journal of Training and Development* 38 (3): 180–197. <https://doi.org/10.1108/EJTD-09-2013-0103>.
- Fenoglio, Enzo, Emre Kazim, Hugo Latapie, and Adriano Koshiyama. 2022. “Tacit Knowledge Elicitation Process for Industry 4.0.” *Discover Artificial Intelligence* 2 (1): 1–15. <https://doi.org/10.1007/s44163-022-00020-w>.
- Fereday, Jennifer, and Eimear Muir-Cochrane. 2006. “Demonstrating Rigor Using Thematic Analysis: A Hybrid Approach of Inductive and Deductive Coding and Theme Development.” *International Journal of Qualitative Methods* 5 (1): 80–92. <https://doi.org/10.1177/160940690600500107>.
- Fox, Sarah E., Vera Khovanskaya, Clara Crivellaro, Niloufar Salehi, Lynn Dombrowski, Chinmay Kulkarni, Lilly Irani, and Jodi Forlizzi. 2020. “Worker-Centered Design: Expanding HCI Methods for Supporting Labor.” In *(CHI EA '20)*, 1–8. New York, NY, USA: Association for Computing Machinery. <https://doi.org/10.1145/3334480.3375157>.
- Gao, Tianyu, Adam Fisch, and Danqi Chen. 2021. “Making Pre-Trained Language Models Better Few-Shot Learners.” In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 3816–3830. Kerrville, USA: Association for Computational Linguistics, Online. <https://doi.org/10.18653/v1/2021.acl-long.295>.
- Gerbracht, Marc, Max Krüger, Débora De Castro Leal, Peter Tolmie, and Volker Wulf. 2024. ““We Are No Luddites!” – CSCW, Co-Determination and Digital Transformation in Germany.” *Proceedings of the ACM on Human-Computer Interaction* 8:153:1–153:23. <https://doi.org/10.1145/3637430>.
- Greenbaum, Joan, and Morten Kyng. 1991. *Design at Work: Cooperative Design of Computer Systems*. Boca Raton, USA: CRC Press. <https://doi.org/10.1201/9781003063988>.
- Hannola, Lea, Francisco Lacueva-Pérez, Paolo Pretto, Alexander Richter, Marlene Schäfer, and Melanie Steinhüser. 2020. “Assessing the Impact of Socio-Technical Interventions on Shop Floor Work Practices.” *International Journal of Computer Integrated Manufacturing* 33 (6): 550–571. <https://doi.org/10.1080/0951192X.2020.1775296>.
- Harbarth, Lydia, Eva Gößwein, Daniel Bodemer, and Lenka Schnaubert. 2024. “(Over)Trusting AI Recommendations: How System and Person Variables Affect Dimensions of Complacency.” *International Journal of Human-Computer Interaction* 41 (1): 1–20. <https://doi.org/10.1080/10447318.2023.2301250>.
- High, Rob. 2012. “The Era of Cognitive Systems: An Inside Look at IBM Watson and how it Works.” *IBM Corporation, Redbooks* 1: 1–16. <https://www.redbooks.ibm.com/redpapers/pdfs/redp4955.pdf>.
- Hoerner, Lorenz, Markus Schamberger, and Freimut Bodendorf. 2022. “Using Tacit Expert Knowledge to Support Shop-Floor Operators through a Knowledge-Based Assistance System.” *Computer Supported Cooperative Work (CSCW)* 32: 55–91. <https://doi.org/10.1007/s10606-022-09445-4>.
- Hoffmann, Sven, Aparecido Fabiano Pinatti de Carvalho, Darwin Abele, Marcus Schweitzer, Peter Tolmie, and Volker Wulf. 2019. “Cyber-Physical Systems for Knowledge and Expertise Sharing in Manufacturing Contexts: Towards a Model Enabling Design.” *Computer Supported Cooperative Work* 28 (3–4): 469–509. <https://doi.org/10.1007/s10606-019-09355-y>.
- Hoffmann, Sven, Thomas Ludwig, Florian Jasche, Volker Wulf, and David Randall. 2022. “RetrofittAR: Supporting Hardware-Centered Expertise Sharing in Manufacturing Settings through Augmented Reality.” *Computer Supported Cooperative Work (CSCW)* 32 (1): 93–139. <https://doi.org/10.1007/s10606-022-09430-x>.
- Israilidis, John, Evangelia Siachou, and Stephen Kelly. 2020. “Why Organizations Fail to Share Knowledge: An Empirical Investigation and Opportunities for Improvement.” *Information Technology & People* 34 (5): 1513–1539. <https://doi.org/10.1108/ITP-02-2019-0058>.
- Jawahar, Ganesh, Benô Sagot, and Djamé Seddah. 2019. “What Does BERT Learn about the Structure of Language?” In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3651–3657. Florence, Italy: Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-1356>.
- Kernan Freire, Samuel, Sarath Surendranadha Panicker, Santiago Ruiz-Arenas, Zoltán Rusák, and Evangelos Niforatos. 2022. “A Cognitive Assistant for Operators: AI-Powered Knowledge Sharing on Complex Systems.” *IEEE Pervasive Computing* 22 (1): 1–9. <https://doi.org/10.1109/MPRV.2022.3218600>.
- Kiangala, Kahiomba Sonia, and Zenghui Wang. 2024. “An Experimental Hybrid Customized AI and Generative AI Chatbot Human Machine Interface to Improve a Factory Troubleshooting Downtime in the Context of Industry 5.0.” *The International Journal of Advanced Manufacturing Technology* 132: 2715–2733. <https://doi.org/10.1007/s00170-024-13492-0>.
- Kjærup, Maria, Mikael B. Skov, Peter Axel Nielsen, Jesper Kjeldskov, Jens Gerken, and Harald Reiterer. 2021. *Longitudinal Studies in HCI Research: A Review of CHI Publications from 1982–2019*. Cham, Switzerland: Springer. https://doi.org/10.1007/978-3-030-67322-2_2.
- Kotthaus, Christoph, Nico Vitt, Max Krüger, Volkmar Pipek, and Volker Wulf. 2022. “Negotiating Priorities on the Shopfloor: A Design Case Study of Maintainers’ Practices.” *Computer Supported Cooperative Work (CSCW)* 32 (1): 141–210. <https://doi.org/10.1007/s10606-022-09444-5>.
- Lakhani, Karim R., Katja Hutter, S. Healy Pokrywa, and Johann Fuller. 2013. “Open Innovation at Siemens.” *Harvard Business School Case* 613-100. <https://www.hbs.edu/faculty/Pages/item.aspx?num=44999>.
- Lee, Soojin, Xinzhu Yang, Jinhee Kim, and Gukdo Byun. 2022. “Effects of Motivational Climate on Knowledge Hiding: The Mediating Role of Work Alienation.” *Behavioral Sciences* 12 (3): 81, 1–12. <https://doi.org/10.3390/bs12030081>.

- Lewis, Patrick, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, et al. 2020. "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks." In *Proceedings of the 34th International Conference on Neural Information Processing Systems (Vancouver, BC, Canada) (NIPS'20)*. Red Hook, NY, USA: Curran Associates Inc., 16 pages, Article 793.
- Leyer, Michael, Alexander Richter, and Melanie Steinhüser. 2018. "Power to the Workers': Empowering Shop Floor Workers with Worker-Centric Digital Designs." *International Journal of Operations & Production Management* 39 (1): 24–42. <https://doi.org/10.1108/IJOPM-05-2017-0294>.
- Li, Chen, Andreas Kornmaaler Hansen, Dimitrios Chrysostomou, Simon Bogh, and Ole Madsen. 2022. "Bringing a Natural Language-Enabled Virtual Assistant to Industrial Mobile Robots for Learning, Training and Assistance of Manufacturing Tasks." In *2022 IEEE/SICE International Symposium on System Integration (SII 2022) (New York, 2022)*, 238–243. New York, USA: IEEE. <https://doi.org/10.1109/SII52469.2022.9708757>.
- Li, Chen, and Hong Ji Yang. 2021. "Bot-X: An AI-based Virtual Assistant for Intelligent Manufacturing." *Multiagent and Grid Systems* 17 (1): 1–14. <https://doi.org/10.3233/MGS-210340>.
- Liebowitz, J. 2002. "A Look at NASA Goddard Space Flight Center's Knowledge Management Initiatives." *IEEE Software* 19 (3): 40–42. <https://doi.org/10.1109/MS.2002.1003451>.
- Lindsay, Robert K., Bruce G. Buchanan, Edward A. Feigenbaum, and Joshua Lederberg. 1993. "DENDRAL: A Case Study of the First Expert System for Scientific Hypothesis Formation." *Artificial Intelligence* 61 (2): 209–261. [https://doi.org/10.1016/0004-3702\(93\)90068-M](https://doi.org/10.1016/0004-3702(93)90068-M).
- Lista Rossetti, Ana Paula, Guilherme Luz Tortorella, Marina Bouzon, Shang Gao, and Toong Khuan Chan. 2023. "Identifying Industry 4.0 Technologies Enablers for Knowledge Management – a Scoping Review." *The TQM Journal* 36 (1): 340–360. <https://doi.org/10.1108/TQM-05-2022-0173>.
- Ludwig, Thomas, Alexander Boden, and Volkmar Pipek. 2017. "3D Printers as Sociable Technologies: Taking Appropriation Infrastructures to the Internet of Things." *ACM Transactions on Computer-Human Interaction* 24 (2): 28 pages, Article 17–28. <https://doi.org/10.1145/3007205>.
- Maedche, Alexander, Christine Legner, Alexander Benlian, Benedikt Berger, Henner Gimpel, Thomas Hess, Oliver Hinz, Stefan Morana, and Matthias Söllner. 2019. "AI-based Digital Assistants: Opportunities, Threats, and Research Perspectives." *Business & Information Systems Engineering* 61 (4): 535–544. <https://doi.org/10.1007/s12599-019-00600-8>.
- Mantravadi, Soujanya, Andreas Dyroy Jansson, and Charles Moller. 2020. "User-Friendly MES Interfaces: Recommendations for an AI-Based Chatbot Assistance in Industry 4.0 Shop Floors." In *Intelligent Information and Database Systems (ACIIDS 2020), PT II (Cham, 2020)*, edited by N. T. Nguyen, K. Jearanaitanakij, A. Selamat, B. Trawinski, and S. Chittayasothorn, 189–201, Vol. 12034. Cham, Switzerland: Springer International Publishing Ag. https://doi.org/10.1007/978-3-030-42058-1_16.
- Martinez-Maldonado, Roberto, Vanessa Echeverria, Gloria Fernandez-Nieto, Lixiang Yan, Linxuan Zhao, Riordan Alfredo, Xinyu Li, et al. 2023. "Lessons Learnt from a Multimodal Learning Analytics Deployment In-the-Wild." *ACM Transactions on Computer-Human Interaction* 31 (1): 1–41. <https://doi.org/10.1145/3622784>.
- Massey, Anne P., Mitzi M. Montoya-Weiss, and Kent Holcom. 2001. "Re-Engineering the Customer Relationship: Leveraging Knowledge Assets at IBM." *Decision Support Systems* 32 (2): 155–170. [https://doi.org/10.1016/S0167-9236\(01\)00108-7](https://doi.org/10.1016/S0167-9236(01)00108-7).
- Meneweger, Thomas, Daniela Wurhofer, Verena Fuchsberger, and Manfred Tscheligi. 2015. "Working Together with Industrial Robots: Experiencing Robots in a Production Environment." In *2015 24th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, 833–838. New York, USA: IEEE. <https://doi.org/10.1109/ROMAN.2015.7333641>.
- Mercedes-Benz. 2023. Benz Tests Chatgpt in Intelligent Vehicle Production. Mercedes-Benz Group. <https://group-mercedes-benz.com/innovation/digitalisation/industry-4-0/chatgpt-in-vehicle-production.html>.
- Monks, Kathy, Edel Conway, Na Fu, Katie Bailey, Grainne Kelly, and Enda Hannon. 2016. "Enhancing Knowledge Exchange and Combination through HR Practices: Reflexivity as a Translation Process." *Human Resource Management Journal* 26 (3): 304–320. <https://doi.org/10.1111/hrmj.v26.3>.
- Mörke, Frauke. 2022. "Inverted Hierarchies on the Shop Floor: The Organisational Layer of Workarounds for Collaboration in the Metal Industry." *Computer Supported Cooperative Work (CSCW)* 31 (1): 111–147. <https://doi.org/10.1007/s10606-021-09415-2>.
- Mucha, Henrik, Carsten Roecker, Thomas Ludwig, Corinna Ogonowski, Martin Stein, Sebastian Robert, Lukas Galla, Martin Hill, and Volker Wulf. 2018. "The Industrial Internet of Things: New Perspectives on HCI and CSCW within Industry Settings." In *Companion of the 2018 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '18 Companion)*, 393–400. New York, NY, USA: Association for Computing Machinery. <https://doi.org/10.1145/3272973.3273009>.
- Müller, Marvin, Emanuel Alexandi, and Joachim Metternich. 2021. "Digital Shop Floor Management Enhanced by Natural Language Processing." *Procedia CIRP* 96:21–26. <https://doi.org/10.1016/j.procir.2021.01.046>.
- Naqvi, Syed Meesam Raza, Mohammad Ghufra, Safa Meraghni, Christophe Varnier, Jean-Marc Nicod, and Nouredine Zerhouni. 2022. "Human Knowledge Centered Maintenance Decision Support in Digital Twin Environment." *Journal of Manufacturing Systems* 65:528–537. <https://doi.org/10.1016/j.jmsy.2022.10.003>.
- Nonaka, Ikujiro. 2009. "The Knowledge-Creating Company." In *The Economic Impact of Knowledge*, 175–187. London, UK.
- Nonaka, Ikujiro, and Hirotaka Takeuchi. 1995. *The Knowledge-Creating Company: How Japanese Companies*

- Create the Dynamics of Innovation. Oxford, UK: Oxford University Press.
- Olander, Heidi, Mika Vanhala, Pia Hurmelinna-Laukkanen, and Kirsimarja Blomqvist. 2016. "Preserving Prerequisites for Innovation: Employee-Related Knowledge Protection and Organizational Trust." *Baltic Journal of Management* 11 (4): 493–515. <https://doi.org/10.1108/BJM-03-2015-0080>.
- Oleksa-Marewska, Karolina. 2020. "Organizational Climate as a Mediating Factor between Occupational Stress and Prosocial Organizational Behaviours in Knowledge-Based Organizations." *European Research Studies Journal* 23 (2): 741–762. <https://doi.org/10.35808/ersj/1619>.
- Oliveira, Mírian, Jaime Teixeira, Carla Maria Marques Curado, and Clécio Falcão Araújo. 2023. "Practices That Mitigate Organizational Knowledge Hiding: A Systematic Literature Review." *European Conference on Knowledge Management* 24 (2). <https://doi.org/10.34190/eckm.24.2.1579>.
- OpenAI. 2023. *GPT-4 Technical Report*. arXiv:2303.08774 [cs.CL].
- Oruç, Orçun. 2020. "A Semantic Question Answering through Heterogeneous Data Source in the Domain of Smart Factory." *International Journal on Natural Language Computing (IJNLC)* 9 (4): 33–52. <https://doi.org/10.5121/ijnlc.2020.9403>.
- Pavlov, Dmitry, Igor Sosnovsky, Vyacheslav Dimitrov, Vasilii Melentyev, and Dmitry Korzun. 2020. "Case Study of Using Virtual and Augmented Reality in Industrial System Monitoring." In *2020 26th Conference of Open Innovations Association (FRUCT)*, 367–375. New York, USA: IEEE. <https://doi.org/10.23919/FRUCT48808.2020.9087410>.
- Probst, Gilbert J. B., Steffen Raub, and Kai Romhardt. 2000. *Managing Knowledge: Building Blocks for Success*. Chichester, UK: Wiley.
- Ras, Eric, Fridolin Wild, Christoph Stahl, and Alexandre Baudet. 2017. "Bridging the Skills Gap of Workers in Industry 4.0 by Human Performance Augmentation Tools: Challenges and Roadmap." In *Proceedings of the 10th International Conference on Pervasive Technologies Related to Assistive Environments (PETRA '17)*, 428–432. New York, NY, USA: Association for Computing Machinery. <https://doi.org/10.1145/3056540.3076192>.
- Reis, Arsénio, João Barroso, Arlindo Santos, Paulo Rodrigues, and Rodrigo Pereira. 2022. "Virtual Assistance in the Context of the Industry 4.0: A Case Study at Continental Advanced Antenna." In *Information Systems and Technologies (Cham, 2022)*, edited by Alvaro Rocha, Hojjat Adeli, Gintautas Dzemyda, and Fernando Moreira, 651–662. Cham, Switzerland: Springer International Publishing. https://doi.org/10.1007/978-3-031-04826-5_64.
- Roeein, Donya, Devis Bianchini, Francesco Leotta, Massimo Mecella, Paolo Paolini, and Barbara Pernici. 2020. "Chatting About Processes in Digital Factories: A Model-Based Approach." In *Enterprise, Business-Process and Information Systems Modeling (Cham, 2020)*, edited by Selmin Nurcan, Iris Reinhardt-Berger, Pnina Soffer, and Jelena Zdravkovic, 70–84. Cham, Switzerland: Springer International Publishing. https://doi.org/10.1007/978-3-030-49418-6_5.
- Schniederjans, Dara G., Carla Curado, and Mehrnaz Khalajhedayati. 2020. "Supply Chain Digitisation Trends: An Integration of Knowledge Management." *International Journal of Production Economics* 220:107439. <https://doi.org/10.1016/j.ijpe.2019.07.012>.
- Sexton, Thurston, Michael P. Brundage, Michael Hoffman, and K. C. Morris. 2017. "Hybrid Datafication of Maintenance Logs from AI-Assisted Human Tags." In *2017 IEEE International Conference on Big Data (Big Data)*, 1769–1777. New York, USA: IEEE. <https://doi.org/10.1109/BigData.2017.8258120>.
- Shneiderman, Ben. 2022. *Human-Centered AI*. Oxford, UK: Oxford University Press.
- Shortliffe, Edward H., Randall Davis, Stanton G. Axline, Bruce G. Buchanan, C. Cordell Green, and Stanley N. Cohen. 1975. "Computer-Based Consultations in Clinical Therapeutics: Explanation and Rule Acquisition Capabilities of the MYCIN System." *Computers and Biomedical Research* 8 (4): 303–320. [https://doi.org/10.1016/0010-4809\(75\)90009-9](https://doi.org/10.1016/0010-4809(75)90009-9).
- Sijbom, Roy B. L., Ellis S. Emanuel, Jessie Koen, Matthijs Baas, and Leander De Schutter. 2025. "Daily Knowledge Sharing at Work: The Role of Daily Knowledge Sharing Expectations, Learning Goal Orientation and Task Interdependence." *European Journal of Work and Organizational Psychology* 34 (2): 298–314. <https://doi.org/10.1080/1359432X.2025.2458343>.
- Strickland, Eliza. 2019. "IBM Watson, Heal Thyself: How IBM Overpromised and Underdelivered on AI Health Care." *IEEE Spectrum* 56 (4): 24–31. <https://doi.org/10.1109/MSPEC.2019.8678513>.
- Tan, Hai Chen, Pat Carrillo, Chimay Anumba, John M. Kamara, Dino Bouchlaghem, and Chika Udejaja. 2006. "Live Capture and Reuse of Project Knowledge in Construction Organisations." *Knowledge Management Research & Practice* 4 (2): 149–161. <https://doi.org/10.1057/palgrave.kmrp.8500097>.
- Tao, Wenjin, Ze-Hao Lai, Ming C. Leu, Zhaozheng Yin, and Ruwen Qin. 2019. "A Self-Aware and Active-Guiding Training & Assistant System for Worker-Centered Intelligent Manufacturing." *Manufacturing Letters* 21:45–49. <https://doi.org/10.1016/j.mfglet.2019.08.003>.
- Trappey, Amy J. C., Charles V. Trappey, Min-Hua Chao, and Chun-Ting Wu. 2022. "VR-enabled Engineering Consultation Chatbot for Integrated and Intelligent Manufacturing Services." *Journal of Industrial Information Integration* 26:100331. <https://doi.org/10.1016/j.jii.2022.100331>.
- van Greunen, Conrad, Elmarie Venter, and Gary Sharp. 2021. "The Influence of Relationship and Task Conflict on the Knowledge-Sharing Intention in Knowledge-Intensive Organisations." *South African Journal of Business Management* 52 (1): Article 2166, 1–9. <https://doi.org/10.4102/sajbm.v52i1.2166>.
- Vasconcelos, Helena, Matthew Jörke, Madeleine Grundle-McLaughlin, Tobias Gerstenberg, Michael S. Bernstein, and Ranjay Krishna. 2023. "Explanations Can Reduce Overreliance on AI Systems during Decision-Making." *Proceedings of the ACM on Human-Computer Interaction* 7 (CSCW1): 38 pages, Article 129–38. <https://doi.org/10.1145/3579605>.

- Vorm, E. S., and David J. Y. Combs. 2022. "Integrating Transparency, Trust, and Acceptance: The Intelligent Systems Technology Acceptance Model (ISTAM)." *International Journal of Human-Computer Interaction* 38 (18): 1828–1845. <https://doi.org/10.1080/10447318.2022.2070107>.
- Wagner, Christian. 2006. "Breaking the Knowledge Acquisition Bottleneck through Conversational Knowledge Management." *Information Resources Management Journal (IRMJ)* 19 (1): 70–83. <https://doi.org/10.4018/irmj.2006010104>.
- Wang, Hongwei, Fuzheng Zhang, Miao Zhao, Wenjie Li, Xing Xie, and Minyi Guo. 2019. "Multi-Task Feature Learning for Knowledge Graph Enhanced Recommendation." In *The World Wide Web Conference*, 2000–2010. San Francisco, CA: Association for Computing Machinery. <https://doi.org/10.1145/3308558.3313411>.
- Wei, Jason, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, et al. 2022a. *Emergent Abilities of Large Language Models*. arXiv:2206.07682 [cs.CL].
- Wei, Jason, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc V. Le, and Denny Zhou. 2022. "Chain-of-Thought Prompting Elicits Reasoning in Large Language Models." *Advances in Neural Information Processing Systems* 35:24824–24837.
- Weinert, Tim, Matthias Billert, Marian Thiel de Gafenco, Andreas Janson, and Jan Marco Leimeister. 2022. "Designing a Co-Creation System for the Development of Work-Process-Related Learning Material in Manufacturing." *Computer Supported Cooperative Work (CSCW)* 32 (1): 5–53. <https://doi.org/10.1007/s10606-021-09420-5>.
- Wellsandt, Stefan, 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* 53:382–390. <https://doi.org/10.1016/j.arcontrol.2022.04.001>.
- Wieland, Matthias, Pascal Hirmer, Frank Steimle, Christoph Groeger, Bernhard Mitschang, Eike Rehder, Dominik Lucke, Omar Abdul Rahman, and Thomas Bauernhansl. 2016. "Towards a Rule-Based Manufacturing Integration Assistant." In *Factories of the Future in the Digital Environment (Amsterdam, 2016)*, edited by E. Westkamper and T. Bauernhansl, 213–218, Vol. 57. Amsterdam, The Netherlands: Elsevier Science Bv. <https://doi.org/10.1016/j.procir.2016.11.037>.
- Wurhofer, Daniela, Thomas Meneweger, Verena Fuchsberger, and Manfred Tscheligi. 2018. "Reflections on Operators' and Maintenance Engineers' Experiences of Smart Factories." In *Proceedings of the 2018 ACM International Conference on Supporting Group Work (GROUP '18)*, 284–296. New York, NY, USA: Association for Computing Machinery. <https://doi.org/10.1145/3148330.3148349>.
- Xia, Yuchen, Manthan Shenoy, Nasser Jazdi, and Michael Weyrich. 2023. *Towards Autonomous System: Flexible Modular Production System Enhanced with Large Language Model Agents*. arXiv:2304.14721 [cs.RO].
- Xu, Xun, Yuqian Lu, Birgit Vogel-Heuser, and Lihui Wang. 2021. "Industry 4.0 and Industry 5.0 – Inception, Conception and Perception." *Journal of Manufacturing Systems* 61:530–535. <https://doi.org/10.1016/j.jmsy.2021.10.006>.
- Yang, Chi-Lan, Chien Wen (Tina) Yuan, and Hao-Chuan Wang. 2019. "When Knowledge Network Is Social Network: Understanding Collaborative Knowledge Transfer in Workplace." *Proceedings of the ACM on Human-Computer Interaction* 3 (CSCW): 1–23. <https://doi.org/10.1145/3359266>.
- Zhang, Rui, Nathan J. McNeese, Guo Freeman, and Geoff Musick. 2021. "'An Ideal Human': Expectations of AI Teammates in Human-AI Teaming." *Proceedings of the ACM on Human-Computer Interaction* 4 (CSCW3): 25 pages, Article 246–25. <https://doi.org/10.1145/3432945>.
- Zhang, Yunfeng, Q. Vera Liao, and Rachel K. E. Bellamy. 2020. "Effect of Confidence and Explanation on Accuracy and Trust Calibration in AI-assisted Decision Making." In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 295–305. New York, NY: Association for Computing Machinery. <https://api.semanticscholar.org/CorpusID:210023849>.
- Zhao, Wayne Xin, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, et al. 2023. *A Survey of Large Language Models*. arXiv:2303.18223 [cs.CL].