



Delft University of Technology

Tacit Knowledge Elicitation for Shop-floor Workers with an Intelligent Assistant

Kernan Freire, Samuel; Wang, Chaofan; Ruiz-Arenas, Santiago; Niforatos, Evangelos

DOI

[10.1145/3544549.3585755](https://doi.org/10.1145/3544549.3585755)

Publication date

2023

Document Version

Final published version

Published in

CHI 2023 - Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems

Citation (APA)

Kernan Freire, S., Wang, C., Ruiz-Arenas, S., & Niforatos, E. (2023). Tacit Knowledge Elicitation for Shop-floor Workers with an Intelligent Assistant. In *CHI 2023 - Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems* Article 266 (Conference on Human Factors in Computing Systems - Proceedings). ACM. <https://doi.org/10.1145/3544549.3585755>

Important note

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

Copyright

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.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Tacit Knowledge Elicitation for Shop-floor Workers with an Intelligent Assistant

S. Kernan Freire
Delft University of Technology
Delft, The Netherlands
s.kernanfreire@tudelft.nl

S. Ruiz-Arenas
Universidad EAFIT
Medellin, Colombia
sruizare@eafit.edu.co

C. Wang
Delft University of Technology
Delft, The Netherlands
c.wang-16@tudelft.nl

E. Niforatos
Delft University of Technology
Delft, Netherlands
e.niforatos@tudelft.nl

ABSTRACT

Many industries face the challenge of capturing workers' knowledge to share it, particularly tacit knowledge. The operation of complex systems such as a manufacturing line is knowledge-intensive. Considering this knowledge's breadth and dynamic nature, existing knowledge-sharing solutions are inefficient and resource intensive. Conversational user interfaces are an efficient way to convey information that mimics how humans share knowledge; however, we know little about how to design them specifically for knowledge sharing, especially regarding tacit knowledge. In this work, we present an intelligent assistant that we have developed to support the elicitation of tacit knowledge from workers through systematic reflection. The system can interact with workers by voice or text and generate visualizations of shop floor data to support reflective prompts.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in HCI**; **Interactive systems and tools**; **Natural language interfaces**.

KEYWORDS

intelligent assistant, chatbots, tacit knowledge, systematic reflection, industry 5.0, human-centered AI, knowledge sharing

ACM Reference Format:

S. Kernan Freire, C. Wang, S. Ruiz-Arenas, and E. Niforatos. 2023. Tacit Knowledge Elicitation for Shop-floor Workers with an Intelligent Assistant. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23)*, April 23–28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3544549.3585755>

1 INTRODUCTION

Effective knowledge management can significantly impact firm performance [12]. Tacit knowledge, which can be classified into cognitive and technical dimensions, plays a crucial role in this

process. Whereas the technical dimension is the people's know-how (i.e., technical skills), the cognitive dimension includes the mental models, beliefs, and values that influence how people perceive the world [28], and the cognitive dimension resides in memory/cognitive processes; It is associated with hunches and intuition [38]. While it is relatively straightforward to describe explicit knowledge, tacit knowledge is difficult to express or codify and is often acquired through experience and practice [32]. Tacit knowledge is not homogeneous; It has explicit components that can be articulated and codified [15]. Nonaka and Takeuchi [32] make a distinction between expressible tacit knowledge and inexpressible tacit knowledge. The expressible part of the know-how includes recipes and formulas, rules of thumb, and tricks of the trade, among others [28]. The inexpressible part is associated with movement skills and physical experiences that someone cannot easily explain or replicate based on a verbal description [28]. On the shop floor, tacit knowledge is critical for the efficient and effective operation of manufacturing processes [18, 31]. This knowledge includes the skills, routines, and practices workers use to perform their tasks and an understanding of the equipment and their workflows. Studies in the manufacturing industry have shown that sharing codified tacit knowledge improves task performance (e.g., [14]).

One of the main difficulties associated with the transfer of tacit knowledge is that people are unaware that they have it; Identifying relevant knowledge is a significant challenge. Therefore, the ability to convert tacit knowledge into explicit knowledge that can be shared is precious for a company [32]. Tacit knowledge is often shared informally among workers and is passed down through generations of employees. An immense volume of research has explored how to support this natural process (e.g., [31] as well as methods to codify the knowledge so that it can be shared at scale and asynchronously (e.g., [14]). These methods to codify tacit knowledge typically involve a skilled analyst to perform interviews, observations, or thought mapping. To reduce the burden on workers and analysts, recent work has explored using assistance systems to support this process [17]; however, they still require a human analyst and significant time investments.

Intelligent assistants (IA), a type of assistance system that implements artificial intelligence to increase human task performance at work [25], are becoming increasingly popular in the home and at businesses. Whereas the most iconic IAs are consumer-facing (e.g., Alexa, Google Assistant, or customer service chatbots), they

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

CHI EA '23, April 23–28, 2023, Hamburg, Germany

© 2023 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9422-2/23/04.

<https://doi.org/10.1145/3544549.3585755>

are increasingly being used in a professional setting (e.g., medical decision support or cognitive assistance on the shop floor (e.g., [21]). Some of the benefits that (voice-enabled) IAs provide include hand and gaze-free interaction, quickly accessing information, processing a large volume of information from their user (as natural language), and can help to visualize and analyze data in real-time. Furthermore, IAs can represent and process domain knowledge and learn new knowledge through machine learning [25].

Many of the techniques for tacit knowledge sharing involve some form of reflection [17, 38]. Reflection helps to elicit tacit skills, enabling knowledge externalization [3]. Systematic reflection is also essential for any professional to continuously learn [8]. Indeed, continuous learning is where a company can improve its competitive advantage [24]. As such, we propose using an IA to facilitate short systematic reflections on the shop floor to elicit tacit knowledge and promote continuous learning. An IA can support efficient, systematic reflection by using data from the production line to visualize a worker's actions and system behavior for a specific activity (e.g., a production batch). Then, it can process the worker's reflections using natural language processing (NLP) to store and reuse the knowledge. The only existing system for sharing tacit knowledge on the shop floor requires a human analyst to perform activities that include interviewing and observing shop floor workers [17]. To our knowledge, no system supports systematic reflection on the shop floor for continuous learning and tacit knowledge elicitation. We will focus on the technical dimension of tacit knowledge. More specifically, we aim to capture expressible tacit knowledge by eliciting heuristics, rules of thumb, and processes from experts and workers on the shop floor. Likewise, we aim to contribute to identifying non-expressible tacit knowledge from a technical perspective (i.e., know-what and know-how) by supporting workers to reflect on their work performance. This project stage does not explore the cognitive dimension, including intuition, hunches, and predictions.

In this paper, we focus on IA-supported systematic reflection; however, this builds on several existing capabilities of our IA. Existing capabilities include the following: collect event reports; collect and share production settings; provide standard work instructions; and offer recommendations to issue handling based on acquired knowledge. Previously, we focused on acquiring explicit knowledge directly (e.g., what are the best settings for a product) and used knowledge graph analytics to identify the tacit knowledge. We aim to use systematic reflection to support workers in making their tacit knowledge explicit and helping them learn.

2 BACKGROUND

The importance of worker knowledge on the shop floor has been widely acknowledged [31]. Agile manufacturing is characterized by a flexible production environment that can adapt to customer demand, for example, by using the same production line to produce numerous discrete products, such as different models of cars. However, it is also applied in process manufacturing, such as chemicals. Process manufacturing involves the creation of products by adding ingredients according to a formula or recipe. The IA we developed has been designed in collaboration with a detergent production company that produces detergents in batches according to client orders before reconfiguring the production line for another order.

In this context, shop floor workers must continuously adjust parameters and fix issues to ensure that they produce at high speed while minimizing stoppages. However, as every product is unique, and numerous external factors affect production (e.g., raw material quality), workers need to adjust strategically. As such, their work is highly knowledge-intensive and dynamic. The sheer frequency of issues and adjustments means that any intervention introduced must minimize the cognitive and temporal impact on their work.

Like in other industries, manufacturing faces challenges in sharing knowledge between workers, especially tacit knowledge. Inherently, tacit knowledge is implicit, making it more difficult for individuals and companies to acquire and disseminate it [18]. Although several definitions of tacit knowledge state that it cannot be expressed verbally, Nonaka and Takeuchi [32] suggest that it can be divided into inexpressible and expressible types. We know that tacit knowledge can be converted into explicit knowledge and exists on a continuum [33]. However, storing it in a repository without considering the social context is not sufficient [32].

Researchers have manually acquired tacit knowledge through numerous techniques, such as expert interviews, human-motion capture, video, and concept maps [3, 5, 16]. There are also several verbal techniques to make tacit knowledge explicit, such as reflective practices [36], thinking aloud [9], and collegial verbalization [10]. In contrast to self-reflection and thinking aloud (about one's actions), collegial verbalization involves reflection on the recorded actions of a colleague, which Erlandsson and Jansson [10] claim to make the process more effective. In addition to the potential to make tacit knowledge explicit to share with others, using systematic reflection can facilitate continuous learning [8]. However, the typical way of performing these activities takes time and often requires a second (human) party to facilitate. In general, the process of eliciting, capturing, and sharing tacit knowledge is resource intensive [16]. Even recent attempts to incorporate assistance systems still rely on skilled analysts to perform several tasks (e.g., interviews and observations) [17]. Whereas this resource-intensive process might be worth conducting in an environment where knowledge is relatively static, it is less viable in an environment where knowledge is dynamic (e.g., due to new product introductions or adjustments to an agile production line). Indeed, Ellis et al. [8] observe that further research is needed to determine how systematic reflection can be integrated into a busy work environment.

In recent years, continuous learning in the workplace has gained increasing interest. With rapid technological advances, many researchers have explored how to facilitate learning through reflection using new technologies. For example, Müller et al. [29] used context data from wearables to support learning through self-reflection in a care home and voluntary crisis workers, and Kocielnik et al. [23] created visualizations of people's experiences using physiological sensor data combined with personal calendar data to support self-reflection in an academic and educational setting. In general, less experienced employees can gain valuable insights from the experiences and viewpoints of more experienced colleagues. On the other hand, more experienced workers may find that individual reflection yields better results for them [37]. In the past few years, researchers have explored using conversational agents to support reflection and workplace learning [13, 22, 34, 41]. For example, Kocielnik et al. [22] developed a multimodal (text and

voice) conversational agent for workplace reflection. They noted that using voice interaction instead of a busy chat modality may allow people to take a step back and reflect on their work. Other work has shown that reporting by voice to an IA results in reports with more information, including valuable explanations, than writing on paper [19]. These studies demonstrate that conversational interaction is an exciting technique for supporting self-reflection in the workplace.

While manufacturing automation has tended to replace humans, we now expect intelligent systems to collaborate with human intelligence [30]. Worker assistance systems support manufacturing operators in completing tasks without replacing them, overruling, or posing any danger to the operator [26]. There are different levels of collaboration in decision-making: assisted, verified, and delegated. Delegated and verified decision-making can be performed when data is highly structured, the task can be standardized, and the decision-making process is not a threat to human life or the environment. In contrast, assisted decision-making is implemented when highly fragmented data from various sources is needed, a high level of expertise is required, and critical consequences may result from the decision-making [25]. Our IA falls into the latter category, as the production of detergents uses dangerous chemicals under high pressure. Therefore, it is also essential to consider the quality of the knowledge the IA acquires.

Industrial applications of Intelligent Assistants, similar to Alexa, Google Assistant, or Siri, are an emerging research topic. Several IAs using a conversational interface have emerged in different research communities with other names (e.g., intelligent (personal) assistants, digital assistants, software robots, or chatbots). IAs can have significant benefits, such as central access to information systems and ubiquitous decision support [2, 21, 39]. Intelligent (robot) systems can support decision-making by analyzing shop floor data, identifying production issues, and evaluating operation performance [1, 35]. In general, a growing body of work shows the positive impact that IAs can have on the shop floor, especially when integrated with live shop floor data.

Live data from the shop floor can be used in many situations, such as assisting with machine troubleshooting, adjusting product formulations, or evaluating shift performance. The simplest form of analysis would be to look at raw data; however, extracting more value from the data using additional processes, such as descriptive statistics, machine learning, and visualizations are possible. Data visualization can also be an effective tool for self-reflection, as it allows people to explore and understand their own behavior and thought patterns. Furthermore, it can provide memory cues to facilitate reflection on a busy 8-hour shift. By creating visual representations of data, people can quickly identify patterns and trends in their work, and gain insight into areas that may require improvement. For example, visualizing data on production rates, downtime, and quality control can help workers understand how their actions impact overall performance and identify areas where they can make changes to improve efficiency and reduce costs. Data visualization can also be used to set goals, monitor progress, and measure success. Whereas data visualization has frequently been used to support learning and reflection, we are unaware of any research exploring using it together with an IA on the shop floor.

Implementing an IA for learning and knowledge sharing will face numerous technical and human challenges. For example, workers may intentionally withhold knowledge to avoid being replaced or reduce their perceived value [6]. Human autonomy is another challenge, as intelligent assistants can prevent people from keeping control of their decisions and actions and lead to incompetence in task performance [25]. Intrusiveness and privacy are two of the most reported challenges in implementing intelligent assistants for workers [11]. The matter of knowledge quality is also of great concern for the factories involved, as it is possible for workers to share bad practices. Finally, it will be challenging to implement a self-reflection activity that fits the busy work environment like the shop floor [8].

3 SYSTEMATIC REFLECTION SUPPORTED BY AN INTELLIGENT ASSISTANT

3.1 Development

The development of the IA for systematic reflection was carried out in close collaboration with factory partners; It builds on several existing features such as collecting issue descriptions and providing recommendations as described in previous literature [20, 21]. The development involved conducting semi-structured interviews, collecting requirements from workers and management, observation, user evaluations, and integrating the IA with the factory's IT systems. The factory workers also helped us develop a simulated production environment and validated its accuracy (see 3.5 for more information). Together, we decided to focus on eliciting (tacit) knowledge related to optimizing production parameters and resolving unplanned stoppages as this would greatly impact production. The shift leaders had observed a significant difference in knowledge between workers on the shop floor on this task. This could be attributed to the highly dynamic nature of the knowledge of the optimization process as new products are introduced, machine behavior is dynamic (e.g., as parts wear down), and the formulas and raw ingredient quality can change.

3.2 Factory Context

As Matthew and Sternberg [27] states, reflection is the key to experiential learning. In the hectic environment of a shop floor, workers do not have much time to reflect on their work. In a factory that operates 24/7, shop floor workers typically work in shifts of 8 hours, which is also the case for the detergent factories with which we collaborated. Workers have periods of high-intensity work, punctuated by periods of waiting. During an observational study we conducted across three shifts, we recorded over 100 issues the worker faced and several periods of unplanned downtime (sometimes taking over an hour), which are opportune moments for reflection. Currently, workers must summarize their shift and indicate any problems in a text box for the next shift to read. However, previous research in the manufacturing context has shown that these types of reports are often missing, incomplete, or written in a way that inhibits future use (e.g., inconsistent use of terminology) [7, 40].

3.3 User Interaction

Upon product batch and shift completion, the IA automatically prompts users to start a reflection session (see Figure 1). Alternatively, users can start a session at their discretion, for example, during a lengthy production line stoppage. At first, the IA will summarize the worker's performance relative to previous batches/shifts in natural language. Then, the IA will use the shop floor data to visualize the entire batch (or shift) for the worker (see Figure 2).




Figure 1: A reflective prompt the IA could give at the end of a batch

The visualization includes general performance indicators such as average speed, the number of stops, parameter adjustments, machine state, and product quality indicators. Then, the IA will ask them to reflect on the activity by asking, "*How did you contribute to the performance observed in the activity?*" [8]. Then, the IA will ask the worker to identify some critical moments. If no moments are indicated by the worker, the IA will use a pseudo-algorithmic approach to select some moments itself (e.g., a period of relatively high or low performance). Next, the IA will present visualizations of these moments (see Figure 2) and ask for more detailed reflections, such as "*Why did you do X or decide Y?*" [8].

At both levels, these questions aim to prompt self-reflection that results in specific explanations [8] such that tacit knowledge can become explicit. The IA has been trained to identify references to specific named entities, such as machine components, their states, and worker actions. It can connect any explanation to the associated nodes in a knowledge graph with the contextual information it can automatically acquire from its live data access. This facilitates sharing of acquired knowledge when a similar situation arises.

3.4 Design Principles and Features

As discussed in the background section, several challenges must be overcome to successfully deploy an Intelligent Agent for self-reflection and knowledge sharing in a factory. These challenges include user acceptance; privacy; limits to human memory and self-awareness; the risk of capturing bad practices and; handling dynamically changing knowledge. The aspects related to user acceptance can be broken down into perceived benefits, perceived ease



Figure 2: Visualizing simulated shop floor data enables workers to review parameter adjustments and assess their impact on the production line. The circles indicate different types of user interactions (e.g., changing a parameter) and the blue dots indicate the weights of individual canisters of detergent

of use, integration into daily operations, and user experience [4]. The following paragraphs outline the IA's features and how they address these challenges.

The IA receives real-time data for two main purposes: to support reflection by generating (memory) cues through descriptive statistics and visualizations and to provide context awareness. Context awareness enables the IA to have shorter and more reliable interactions with workers, thus improving its perceived ease of use. For example, it can suggest user prompts by detecting which machines the workers recently stood next to or what product is being produced.

The IA uses natural language processing to extract relevant entities and knowledge (e.g., machine names, solutions, and strategies) from the workers' reflections and populate a Neo4j¹ knowledge graph. This process leverages the capabilities of state-of-the-art Large Language Models (LLM), such as gpt-3.5-turbo², to process unstructured text. However, if the IA cannot extract key information (e.g., which machine component the user referenced), it will ask for it directly. We are also exploring how to use LLMs to generate follow-up reflection prompts to support multi-turn reflection interactions.

Before we introduced the features described in this paper (supporting systematic reflection), the IA could already perform several tasks that workers perceived as beneficial, for example, they could look up standard work instructions and share production settings

¹<https://neo4j.com/>—last accessed March 14, 2023

²<https://platform.openai.com/docs/guides/chat>—last accessed March 14, 2023

and issue solutions. As such, the workers already perceived the IA as potentially beneficial. We expect them to see the value of systematic reflection for tacit knowledge sharing, as they can also benefit from the explicit knowledge that emerges from it. For example, if they face a similar problem in several months but forget how they had previously solved it. However, further evaluation is necessary to assess their acceptance of the features described in this paper.

Currently, the IA supports a user rating system for the knowledge it shares. Once the IA is deployed at the factory, we will explore techniques to check for inconsistencies in its knowledge base or during reflection. Factory knowledge managers, such as expert workers, will check knowledge with a poor rating or inconsistencies. These knowledge managers will also be used to assess the quality of the knowledge shared with the IA before approving it. These measures help maintain the quality and relevance of the knowledge base in a manageable way.

3.5 Evaluation Setup

Before introducing the IA in the wild, we will test it in a simulated production environment. This allows us to develop and evaluate it in a safe, controlled environment. The simulation is based on the production environment in a partner factory. We modeled the behavior of several production machines, including the machine that fills cans with detergent. The machine's behavior depends on the detergent being produced and several other factors, such as the level of the buffer tank, the air bubbles in the pipes, and the machine parameters adjusted by the workers (see Figure 3a) and weights of the produced cans (see Figure 3b). The partners' factories employees have validated and approved the simulation accuracy.

4 RESEARCH PLAN AND DISCUSSION

We plan on conducting user studies in the lab in conjunction with studies in-situ, at the factory. We will compare four conditions, namely, (1) unsupported reflection with free-text entry into a text box (the current situation at the factory), (2) reflection supported by natural language prompts only, (3) reflection supported by natural language prompts and tables (in a chat interface), and finally, (4) reflection supported by natural language prompts and graphical data visualizations. We will conduct a between-groups user study ($n > 40$) in the lab to measure the effect of using the IA for systematic reflection on tacit knowledge sharing, workload, and user experience. Then, we will conduct a within-group study at the factory with a smaller group of participants ($n = 10$). After the user study at the factory, we will organize focus group sessions and semi-structured interviews to explore further improvements and barriers to implementation.

We aim to answer the following research questions (RQ):

RQ1 *Can systematic reflection supported by an IA and data visualization result in tacit knowledge elicitation on the shop floor?*

We think that our interventions will result in tacit knowledge sharing on the shop floor; however, we expect to face knowledge quality issues due to the limitations of NLP compared to a human analyst. We will recruit factory knowledge managers to code the responses to measure the validity and quality of the elicited knowledge. Previous work has shown that systematic reflection is a valid technique

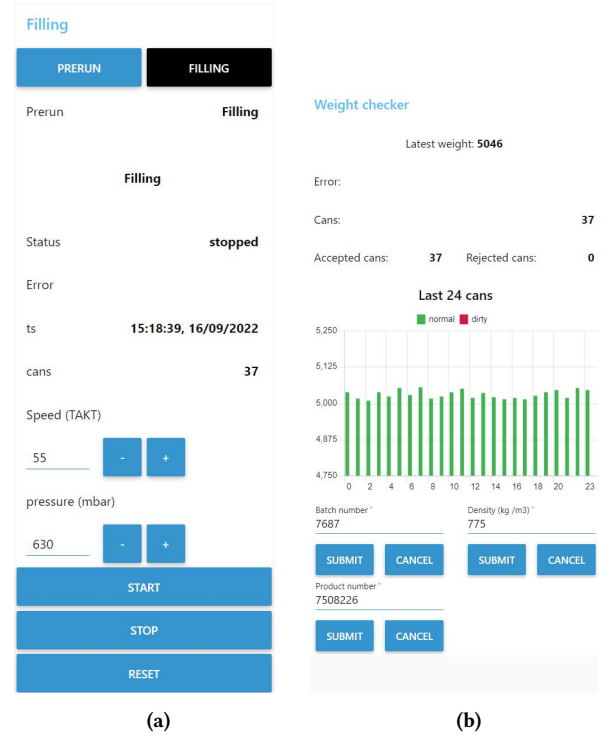


Figure 3: (a) The simulated user interface for the filling machine and (b) the simulated weight checker interface.

for eliciting tacit knowledge [8]. Furthermore, conversation agents can support reflection and learning while on the job [13, 22, 34, 41]; however, they still face reliability issues. Hoerner et al. [17] have shown that assistance systems can support sharing tacit knowledge on the shop floor but still rely largely on human analysts. We expect our knowledge base to be more up-to-date and require less effort to populate as human analysts are unnecessary. However, we do not expect to achieve the same level of knowledge quality.

RQ2 *How do different IA-facilitated systematic reflection techniques affect user acceptance, user experience, and tacit knowledge elicitation on the shop floor?*

The techniques in question are the following: textbox (baseline); natural language prompts only; natural language prompts with text-based data visualization; and natural language prompts with graphical data visualizations. We think that graphical data visualizations will result in more effective self-reflection, as demonstrated by Kocielnik et al. [23], Müller et al. [29] but may require additional interaction steps (i.e., the worker will need to look at another user interface). Ultimately, we will likely face a trade-off, so we want to understand exactly how much value the graphical visualizations bring and at what cost.

RQ3 *What are the barriers to using IA-facilitated systematic reflection on the shop floor?*

One possible challenge is that employees may be resistant to the idea of self-reflection, as it may be seen as intrusive or

unnecessary. In a practical sense, user acceptance might be challenging as it might not instantly provide tangible value to the workers (perceived benefit [4]) and it requires they take time out of their busy schedule. Furthermore, ensuring that the tool is used consistently and effectively across the organization may be difficult. Additionally, there may be concerns about privacy and how the information collected through the tool is used and stored. We believe that involving the shop-floor workers throughout the development and evaluation process will help mitigate some of the above concerns; however, it will remain a sensitive topic as worker knowledge is highly valuable to them and the management.

4.1 Ethical Considerations

Using an Intelligent Assistant (IA) to support systematic reflection on the shop floor raises several ethical considerations. One concern is the privacy and security of the knowledge collected by the IA. It is important to ensure that the knowledge is collected, stored, and used in a way that complies with relevant regulations and laws, such as the General Data Protection Regulation (GDPR) in the European Union. Additionally, it is important to be transparent with employees about how the knowledge is being used and to obtain their informed consent before collecting it.

Another ethical consideration is the impact of the IA on the workforce. It is important to consider how the IA may affect the dynamics and relationships among employees, such as by creating divisions between those who are more willing to use the IA and those who are not. Furthermore, the IA may impose a high cognitive load or distract workers from their other responsibilities when asking them to reflect, especially when initiated at a poor time. Additionally, it is important to consider the potential for bias in the knowledge acquired, analyzed, and shared by the IA. Compared to a human analyst, the IA may be less capable of validating the quality of the acquired knowledge.

The study described in this paper has been approved by a Human Research Ethics Committee.

5 CONCLUSION

In this paper, we have presented the development of an Intelligent Assistant (IA) that supports systematic reflection on the shop floor. The IA is designed to help workers make their tacit knowledge explicit and to assist with learning. The development of the IA was conducted in close collaboration with factory partners and focused on optimizing production parameters. The IA can process highly dynamic knowledge and is designed for reliable and efficient user interaction. However, it is important to consider the ethical implications of introducing an IA in the workplace, including privacy and security, impact on the workers, and the potential for bias.

ACKNOWLEDGMENTS

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 Labor force supported by trustworthy Artificial Intelligence" (Grant agreement 957296).

REFERENCES

- [1] Brunno Abner, Ricardo J. Rabelo, Saulo P. Zambiasi, and David Romero. 2020. Production Management as-a-Service: A Softbot Approach. In *Advances in Production Management Systems. Towards Smart and Digital Manufacturing*, Bojan Lalic, Vidosav Majstorovic, Ugljesa Marjanovic, Gregor von Cieminski, and David Romero (Eds.). IFIP Advances in Information and Communication Technology, Vol. 592. Springer International Publishing, Cham, 19–30. https://doi.org/10.1007/978-3-030-57997-5_13
- [2] André da Silva Barbosa, Felipe Pinheiro Silva, Lucas Rafael dos Santos Crestani, and Rodrigo Bueno Otto. 2018. Virtual assistant to real time training on industrial environment. In *Transdisciplinary Engineering Methods for Social Innovation of Industry 4.0*. IOS Press, 33–42.
- [3] Anitha Chennamaneni and James T. C. Teng. 2011. An Integrated Framework for Effective Tacit Knowledge Transfer. *AMCIS 2011 Proceedings - All Submissions* (Aug. 2011). https://aisel.aisnet.org/amcis2011_submissions/277
- [4] Fred D. Davis. 1989. Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology. *MIS Quarterly* 13, 3 (1989), 319–340. <http://www.jstor.org/stable/249008>
- [5] Cláudio Roberto do Rosário, Liane Mahlmann Kipper, Rejane Frozza, and Bruna Bueno Mariani. 2015. Modeling of tacit knowledge in industry: Simulations on the variables of industrial processes. *Expert Systems with Applications* 42, 3 (Feb. 2015), 1613–1625. <https://doi.org/10.1016/j.eswa.2014.09.023>
- [6] Susanne Durst and Malgorzata Zieba. 2019. Mapping knowledge risks: towards a better understanding of knowledge management. *Knowledge Management Research & Practice* 17, 1 (Jan. 2019), 1–13. <https://doi.org/10.1080/14778238.2018.1538603>
- [7] Brett Edwards, Michael Zatorsky, and Richi Nayak. 2008. Clustering and classification of maintenance logs using text data mining. *Volume 87-Data Mining and Analytics 2008* (2008), 193–199.
- [8] Shmuel Ellis, Bernd Carette, Frederik Anseel, and Filip Lievens. 2014. Systematic Reflection: Implications for Learning From Failures and Successes. *Current Directions in Psychological Science* 23, 1 (2014), 67–72. <https://doi.org/10.1177/0963721413504106>
- [9] K Anders Ericsson and Herbert A Simon. 1984. *Protocol analysis: Verbal reports as data*. the MIT Press.
- [10] M. Erlandsson and A. Jansson. 2007. Collegial verbalisation – a case study on a new method on information acquisition. *Behaviour & Information Technology* 26, 6 (2007), 535–543. <https://doi.org/10.1080/01449290600801934>
- [11] Markus Funk, Tilman Dingler, Jennifer Cooper, and Albrecht Schmidt. 2015. Stop helping me - I'm bored!: why assembly assistance needs to be adaptive. *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2015 ACM International Symposium on Wearable Computers - UbiComp '15* (2015), 1269–1273. <https://doi.org/10.1145/2800835.2807942> Conference Name: the 2015 ACM International Joint Conference.
- [12] Mohamad H Gholami, Mehrdad Nazari Asli, Salman Nazari-Shirkouhi, and Ali Noruzi. 2013. Investigating the influence of knowledge management practices on organizational performance: an empirical study. *Acta Polytechnica Hungarica* 10, 2 (2013), 205–216.
- [13] Ted Grover, Kael Rowan, Jina Suh, Daniel McDuff, and Mary Czerwinski. 2020. Design and Evaluation of Intelligent Agent Prototypes for Assistance with Focus and Productivity at Work. In *Proceedings of the 25th International Conference on Intelligent User Interfaces* (Cagliari, Italy) (IUI '20). Association for Computing Machinery, New York, NY, USA, 390–400. <https://doi.org/10.1145/3377325.3377507>
- [14] Claire Gubbins, Siobhan Corrigan, Thomas N Garavan, Christy O'Connor, Damien Leahy, David Long, and Eamonn Murphy. 2012. Evaluating a tacit knowledge sharing initiative: A case study. *European Journal of Training and Development* (2012). <https://doi.org/10.1108/03090591211263558>
- [15] Jia Hao, Qiang-Fu Zhao, Yan Yan, and Guo-Xin Wang. 2016. A brief introduction to tacit knowledge and the current research topics. In *2016 International Conference on Machine Learning and Cybernetics (ICMLC)*, Vol. 2. 917–921. <https://doi.org/10.1109/ICMLC.2016.7873009> ISSN: 2160-1348.
- [16] Hiroshi Hashimoto, Ikuyo Yoshida, Yoshinori Teramoto, Hideki Tabata, and Chao Han. 2011. Extraction of tacit knowledge as expert engineer's skill based on mixed human sensing. In *2011 Ro-Man*. IEEE, 413–418.
- [17] Lorenz Hoerner, 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)* (Sept. 2022). <https://doi.org/10.1007/s10606-022-09445-4>
- [18] Jeremy Howells. 1996. Tacit knowledge. *Technology analysis & strategic management* 8, 2 (1996), 91–106. <https://doi.org/10.1080/09537329608524237>
- [19] Samuel Kernan Freire, Evangelos Niforatos, Zoltan Rusak, Doris Aschenbrenner, and Alessandro Bozzon. 2022. A Conversational User Interface for Instructional Maintenance Reports. In *Proceedings of the 4th Conference on Conversational User Interfaces* (Glasgow, United Kingdom) (CUI '22). Association for Computing Machinery, New York, NY, USA, Article 5, 6 pages. <https://doi.org/10.1145/3543829.3544516>

- [20] Samuel Kernan Freire, Evangelos Niforatos, Chaofan Wang, Santiago Ruiz-Arenas, Mina Foosherian, Stefan Wellsandt, and Alessandro Bozzon. 2023. Lessons Learnt from Designing and Evaluating CLAICA: A Continuously Learning AI Cognitive Assistant (*IUI '23*). Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3581641.3584042> In press.
- [21] Samuel Kernan Freire, 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* (2022), 1–9. <https://doi.org/10.1109/MPRV.2022.3218600>
- [22] Rafal Kocielnik, Daniel Avrahami, Jennifer Marlow, Di Lu, and Gary Hsieh. 2018. Designing for Workplace Reflection: A Chat and Voice-Based Conversational Agent. In *Proceedings of the 2018 Designing Interactive Systems Conference* (Hong Kong, China) (*DIS '18*). Association for Computing Machinery, New York, NY, USA, 881–894. <https://doi.org/10.1145/3196709.3196784>
- [23] Rafal Kocielnik, Fabrizio Maria Maggi, and Natalia Sidorova. 2013. Enabling self-reflection with LifelogExplorer: Generating simple views from complex data. In *2013 7th International Conference on Pervasive Computing Technologies for Healthcare and Workshops*. 184–191. <https://doi.org/10.4108/icst.pervasivehealth.2013.251934>
- [24] Meng Li and Fei Gao. 2003. Why Nonaka highlights tacit knowledge: a critical review. *Journal of knowledge management* (2003). <https://doi.org/10.1108/13673270310492903>
- [25] Alexander Maedche, Christine Legner, Alexander Benlian, Benedikt Berger, Henner Gimpel, Thomas Hess, Oliver Hinz, Stefan Morana, and Matthias Söllner. 2019. AI-Based Digital Assistants. *Business & Information Systems Engineering* 61, 4 (Aug. 2019), 535–544. <https://doi.org/10.1007/s12599-019-00600-8>
- [26] Benedikt G. Mark, Erwin Rauch, and Dominik T. Matt. 2021. Worker assistance systems in manufacturing: A review of the state of the art and future directions. *Journal of Manufacturing Systems* 59 (2021), 228–250. <https://doi.org/10.1016/j.jmsy.2021.02.017>
- [27] Cynthia T. Matthew and Robert J. Sternberg. 2009. Developing experience-based (tacit) knowledge through reflection. *Learning and Individual Differences* 19, 4 (Dec. 2009), 530–540. <https://doi.org/10.1016/j.lindif.2009.07.001>
- [28] Haradhan Mohajan. 2016. Sharing of tacit knowledge in organizations: a review. *American Journal of Computer Science and Engineering* 3, 2 (2016), 6–19.
- [29] Lars Müller, Monica Divitini, Simone Mora, Verónica Rivera-Pelayo, and Wilhelm Stork. 2015. Context Becomes Content: Sensor Data for Computer-Supported Reflective Learning. *IEEE Transactions on Learning Technologies* 8, 1 (2015), 111–123. <https://doi.org/10.1109/TLT.2014.2377732>
- [30] Saeid Nahavandi. 2019. Industry 5.0—A human-centric solution. *Sustainability* 11, 16 (2019), 4371. <https://doi.org/10.3390/su11164371>
- [31] Davi Nakano, Jorge Muniz, and Edgard Dias Batista. 2013. Engaging environments: tacit knowledge sharing on the shop floor. *Journal of Knowledge Management* (2013). <https://doi.org/10.1108/13673271311315222>
- [32] Ikujiro Nonaka and Hirotaka Takeuchi. 1995. The knowledge-creating company: How Japanese companies create the dynamics of innovation. *New York, NY* (1995).
- [33] Ikujiro Nonaka and Georg Von Krogh. 2009. Perspective—Tacit knowledge and knowledge conversion: Controversy and advancement in organizational knowledge creation theory. *Organization science* 20, 3 (2009), 635–652.
- [34] Eleftherios Papachristos, Dorte P. Meldgaard, Iben R. Thomsen, and Mikael B. Skov. 2021. ReflectPal: Exploring Self-Reflection on Collaborative Activities Using Voice Assistants. In *Human-Computer Interaction – INTERACT 2021*, Carmelo Ardito, Rosa Lanzilotti, Alessio Malizia, Helen Petrie, Antonio Piccinno, Giuseppe Desolda, and Kori Inkpen (Eds.). Springer International Publishing, Cham, 187–208.
- [35] Ricardo J. Rabelo, Saulo Popov Zambiasi, and David Romero. 2019. Collaborative Softbots: Enhancing Operational Excellence in Systems of Cyber-Physical Systems. In *Collaborative Networks and Digital Transformation*, Luis M. Camarinha-Matos, Hamideh Afsarmanesh, and Dario Antonelli (Eds.). IFIP Advances in Information and Communication Technology, Vol. 568. SPRINGER NATURE, 55–68. https://doi.org/10.1007/978-3-030-28464-0_16
- [36] Eduardo Ravanal Moreno, Valeria M Cabello, Francisco López-Cortés, and Elias Amórtegui Cedeño. 2021. The reflective practice as a tool for making tacit knowledge explicit. *Reflective Practice* 22, 4 (2021), 515–530. <https://doi.org/10.1080/14623943.2021.1930527>
- [37] Bettina Renner, Gudrun Wesiak, Viktoria Pammer-Schindler, Michael Prilla, Lars Müller, Dalia Morosini, Simone Mora, Nils Faltin, and Ulrike Cress. 2020. Computer-supported reflective learning: how apps can foster reflection at work. *Behaviour & Information Technology* 39, 2 (2020), 167–187. <https://doi.org/10.1080/0144929X.2019.1595726>
- [38] Pallavi Satsangi. 2019. Automation of Tacit Knowledge Using Machine Learning. In *2019 6th International Conference on Soft Computing & Machine Intelligence (ISCMI)*. 35–39. <https://doi.org/10.1109/ISCMI47871.2019.9004290> ISSN: 2640-0146.
- [39] Manex Serras, Laura García-Sardiña, Bruno Simões, Hugo Álvarez, and Jon Arambarri. 2020. Dialogue enhanced extended reality: Interactive system for the operator 4.0. *Applied Sciences* 10, 11 (2020), 3960. <https://doi.org/10.3390/app10113960>
- [40] Thurston Sexton, 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. <https://doi.org/10.1109/BigData.2017.8258120>
- [41] Irmtraud Wolfbauer, Viktoria Pammer-Schindler, Katharina Maitz, and Carolyn P. Rosé. 2022. A Script for Conversational Reflection Guidance: A Field Study on Developing Reflection Competence With Apprentices. *IEEE Transactions on Learning Technologies* 15, 5 (2022), 554–566. <https://doi.org/10.1109/TLT.2022.3207226>