

Actions and their Consequences? Implicit Interactions with Workplace Knowledge Bases

Siân Lindley

Microsoft Research, sianl@microsoft.com

Denise Wilkins

Independent researcher, denisewilkins2021@gmail.com

Britta Burlin

Microsoft Research, brburlin@microsoft.com

Innovations in machine learning are enabling organisational knowledge bases to be automatically generated from employees' activities, the results of which can then be presented to workers via the software applications they commonly use. The potential for these systems to shift the ways in which knowledge is produced and shared raises questions regarding what types of knowledge might be inferred from employees' practices, how these can be used to support work, and what the broader ramifications of this might be. This paper draws on findings from two studies to offer an initial exploration of these topics. The research described investigated workplace (i) collaborative actions and (ii) knowledge actions, to explore how they might (i) inform automatically generated knowledge bases, and (ii) find support through the design of intelligent systems. We draw on the literature on implicit interactions in considering next steps.

CCS CONCEPTS • Human-centered computing ~ Empirical studies in HCI • Information systems ~ Enterprise information systems

Additional Keywords and Phrases: machine learning, artificial intelligence, collaboration, information, process, knowing, practice

ACM Reference Format:

Siân Lindley, Denise Wilkins, and Britta Burlin. 2021. In *Proceedings of AutomationXP'21: Workshop on Automation Experience at the Workplace*. In conjunction with CHI'21, May 07, 2021.

1 INTRODUCTION

Innovations in machine learning are enabling knowledge bases (KBs) to be automatically generated from content produced within large organisations, and then presented to workers via the software applications they commonly use. For instance, the recently released *Microsoft Viva Topics* [8] recognises common topics within an organisation, creates 'topic pages' and 'topic cards', and highlights these across Microsoft 365. The increasing capabilities of such systems raise questions about what can be inferred from the activities of employees, and how – in turn – the information that is mined can then be used to support working practices.

Human action is inherently bound up with knowing how to get things done in organisational work. Orlikowski [9] emphasizes that knowing and knowledgeable ability are continuously developed and demonstrated through

Workshop proceedings *Automation Experience at the Workplace*

In conjunction with CHI'21, May 7th, 2021, Yokohama, Japan

Copyright © 2021 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Website: <http://everyday-automation.tech-experience.at>

workers' actions, which are themselves contextual. Thus, knowing is the capacity to perform useful practices, given the unique circumstances of the current situation. Drawing on this position, we present findings from two studies of how actions taken as part of work could (i) inform an automatically generated KB, and (ii) find support through an intelligent system. In Study 1, we asked participants to focus on collaboration; the aim was to identify actions that workers take in relation to others (hereafter 'collaborative actions'). In Study 2, we asked participants to focus on knowledge; the aim was to identify actions taken in relation to important types of knowledge in the workplace (hereafter 'knowledge actions'). It is worth noting that, in Orlikowski's view, knowing cannot be captured; it is inseparable from its constituting practice. However, we posit that information mined from workers' activities could play a role in the performance of useful practices, either by supporting workers' own recurring activities, or by helping others develop the ability to perform competent work. In this short paper, we analyse the collaborative and knowledge actions workers report to highlight possible scenarios for design, and draw out related challenges. We draw on the concept of implicit interactions (e.g., [7] [10]) in considering how to address these.

2 STUDY 1: COLLABORATIVE ACTIONS

Study 1 was a diary study in which participants recorded and then discussed collaborative actions performed as part of work. While the aim was to identify actions that could receive support from intelligent systems, we did not limit data collection to technology-mediated actions, as we were also interested in those accomplished without digital artefacts that may nevertheless raise implications for research or design.

19 participants, who varied in occupation (10 managers/directors; 7 employees, 2 sole traders), age and gender, completed the study. One participant withdrew after the first interview and is not included in the analysis.

The study began with a telephone interview to explain the goals and methods of the study, ascertain consent, and learn about the participant's work. Questions included: What is a typical day like? Who do you work with? What does collaboration look like in your role? How is collaboration managed? What tools do you use as part of collaboration? Participants then completed a diary over four working days and one day off (even if nothing happened on the day off). They were asked to record a few collaborative actions per day, including those that involved no technology or that seemed brief or non-intensive. Finally, participants were interviewed for a second time, in-person, for up to 90 minutes. They talked through the actions they had recorded and, for each action, the interviewer recorded a short description on a sticky note. The sticky notes were organized into groups of similar activities by the participant after all diarized actions had been discussed. Participants were then thanked and given a gift voucher. Interviews were audio-recorded and photos of the sticky note groupings were taken.

3 STUDY 2: KNOWLEDGE ACTIONS

Study 2 was a design workshop in which information workers discussed actions that automatically generated KBs should support. The aim was to understand the everyday actions on knowledge that a KB should enable, and the context in which these actions are performed.

12 participants in two workshops were first shown a video of a fictional AI system that extracted descriptive and procedural information from work and surfaced that knowledge to others to guide next steps. To prime participants to think about knowledge, they were then asked to list important types of knowledge that they use every day for work. Next, participants were asked to generate the key actions they take on knowledge as part of work that should be supported by an ML system. After this, participants were asked to make a flow diagram

to describe a time when they worked collaboratively to achieve a shared objective, highlighting important actions and pieces of knowledge. Finally, in one of the workshops, participants were asked to perform an evaluation exercise, explicitly comparing how the in-context activities described during the flow diagram task fitted, or extended, the collaborative actions that were generated through Study 1. The sessions were audio-recorded, and video recordings and photographs were captured at various points throughout.

4 ANALYSIS AND FINDINGS

For each of the studies, activities generated and described by participants were clustered. The clusters were then organized to produce a single framework across both studies. This framework has three high-level themes: actions that focus on content, actions that focus on process, and actions that focus on a worker’s individual growth, as shown in Table 1.

Table 1: Actions organized into three categories: Actions around *content* include those relating to its generation and maintenance; information seeking; and information verification. Actions around *process* include those relating to learning from other processes, setting up activities and getting buy-in from team members; the management and adaptation of ongoing processes; and action take to gain approval to move an activity forward. Actions around a *worker’s individual growth* relate to building knowledge and gauging expectations relating to one’s role.

Content	Process	Growth
Create, maintain, transform	Scope, set up, build consensus	Grow, teach
Find, acquire, recommend	Manage, adapt, coordinate	Gauge expectations
Assess, verify	Gain approval	

In considering (i) how these actions could inform the development of machine learned KBs, and (ii) how they might receive support from intelligent systems, we developed the framework presented in Figure 1. We applied this framework to contemplate how actions relating to content or to process can result in implicit or explicit contributions to machine-learned systems. We then generated scenarios that might receive support from an intelligent system for each set of actions, and identified related considerations for design. A concise overview of these results is presented in Table 2. A reflection on how organisational KBs might play a role in supporting a worker’s individual growth is the focus of a separate paper, which is currently in preparation.

5 DISCUSSION

Our analysis of collaborative and knowledge actions taken as part of work highlights how employees might make explicit and implicit contributions to organisational KBs, and conversely, how those KBs could support employees as they work. However, it also highlights considerations for design that speak to numerous issues. In this discussion, we draw on the literature on implicit interactions to engage with some of these. Serim & Jacucci [10] define implicit interactions as those “in which the appropriateness of a system response to the user input (i.e., an effect) does not rely on the user having conducted the input to intentionally achieve it”. Their analysis reveals ‘unintentionality’ and ‘unawareness’ to be two ways in which interactions may be implicit, both of which have relevance here. For instance, our findings highlight how the automatic building of organisational KBs could be informed by the implicit contributions of workers, as well as from the implicit feedback they provide by, for instance, selecting one piece of content over another. It seems quite possible that these outcomes would be both unintentional and outside of the user’s awareness, at least initially.

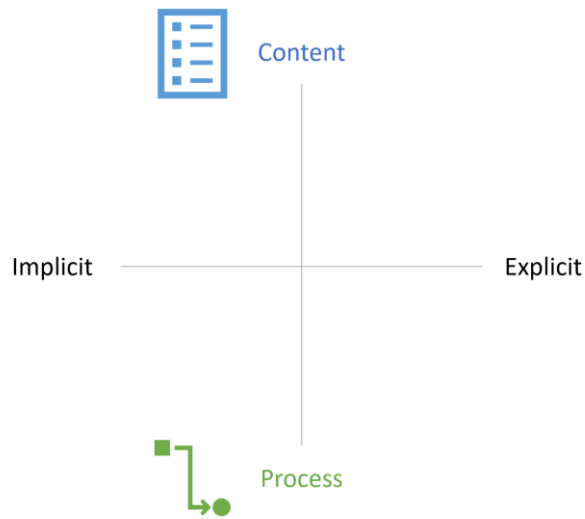


Figure 1: Framework that draws a distinction between worker contributions to a machine learned KB as implicit (i.e., made through work) or explicit (i.e., made knowingly and deliberately), and as relating to content or process.

A lack of awareness over implicit input may be acceptable for systems that have minimal consequences for users, such as the ordering of world-wide web search engine results. However, implicit interactions with organisational KBs have the potential to have more significant effects for workers. As the functionality of such systems begins to mature, the technology could introduce radical shifts in the ways in which knowledge is generated, shared, and consumed within workplaces, which could in turn affect the ways in which employees contribute to, or are *seen* to contribute to, their organisations. Indeed, as workers grow to understand the potential for their work to be made available to others via organisational KBs, it seems possible that they may begin to change their working practices accordingly. Dix [3] has observed that the same interaction can be understood as incidental, expected, or intended, depending on the user's awareness and understanding of the consequences of their actions. Building on this, Serim & Jacucci note that users may avoid certain actions to prevent what they perceive as the unwanted effects of implicit interactions, or they may reformulate their goals in relation to these. In their review, Serim & Jacucci cite research showing how users abstain from reading email to avoid sending read receipts [5], play music not to listen to it but to shape their social media profiles [11], and interact with smart thermostats to mitigate limitations in the technology's modelling of their behaviour [12]. It seems probable then, that employees may seek to retain control over how they, or their work, is presented via organisational KBs by, for instance, deliberately keeping some content private while making other content available via cloud storage services that they expect to be mined by ML systems. Such behaviours resonate with the 'irony of automation' [2], whereby the introduction of automation can necessitate significant changes to action.

This type of agency is dependent on users having some understanding of how organisational KBs are produced and presented to workers. However, workplace systems that are both automatic and pervasive raise challenges in communicating to users what the consequences of their actions are. Janssen et al. [6] argue that the use of automated systems by non-professional users necessitates a deep consideration of how to foster

Table 2: Possible contributions to a KB via content and process-related actions, potential scenarios, and considerations.

	Explicit	Implicit	Ways of supporting work	Considerations
Create, maintain, transform	Add content to KB, update it, edit its status	Content that is created / worked with is included in the KB	Support documentation and sharing	If content is configured out of source material and presented differently via the KB, how is context communicated, authorship acknowledged and accountability determined?
Find, acquire, recommend	Search the KB for content / share via the KB	Content that is shared between colleagues is added to the KB; info about who knows what informs the KB	Help workers find information or the people who can provide it	What are the risks that search results or recommendations will introduce or reflect filter bubbles and biases, and affect workers' opportunities to have impact?
Assess, verify	Verify the accuracy of content within the KB	Verification done in the context of social interaction informs the KB	Support cross-checking or highlight the people who can verify content; confirm that something is correct	When can AI be assigned the authority to verify accuracy and when is a human-in-the-loop needed?
Scope, set up, build consensus	Add info about a process to the KB	Info about teams, timelines, etc. informs the KB	Help workers find organisational policies and learn about processes; generate templates	What are the risks to agility and innovation when processes are replicated and can be followed 'automatically'? Who is the 'permitted' audience for a process?
Manage, adapt, coordinate	Add info about e.g., objectives	Nature of resources, actions, etc. informs the KB	Support management by suggesting actions, giving updates on use of resources; notifying workers regarding policy changes	What does it mean to share know-how via a KB? Are workers possessive over know-how in a way that is different to content?
Gain approval	Sign off content that is entered into the KB	Signoffs done in the context of social interaction informs the KB	Support cross-checking or highlight the people who can verify content; confirm that something is correct	When can AI be assigned the authority to sign-off a proposed action? When is a human-in-the-loop needed?

acceptance in use and, ultimately, trust. Likewise, Fröhlich et al. [4] note that “naïve users” present a unique set of challenges for the design of automated systems. Fröhlich et al. are primarily referring to users who have chosen to use automation technologies outside of work. However, we suggest that both observations could extend to employees who are using, for instance, standard office software and cloud storage technologies, which are deployed by their organisations. Guidelines for the design of AI [1] emphasise responsibility in intelligent systems design, but they tend to do so from the standpoint of intentional use. We see a need to extend these to cover implicit interactions, such that employees using office technologies might gradually build an understanding of how the content they produce and store using workplace technologies is mined by ML systems and made visible in new ways.

In our own research, we are exploring ways of supporting employees in learning how organisational KBs might be shaped by and in turn might shape their work. In this, we take inspiration from Ju & Leifer’s [7] emphasis on the value of designing “courteous” implicit interactions. In particular, we are exploring how systems might do employees the courtesy of helping them understand how their work has contributed to organisational KBs and the practices they support, in the hope of aligning with Serim & Jacucci’s recommendation that implicit interactions should be considered appropriate in retrospect. Our aim is that, by highlighting how implicit interactions have informed organisation KBs, we will enable workers to build mental models of the system that can underpin interactions with it over time and offer opportunities for repair where needed. In doing so, we hope to design user experiences that enable agency and systems that are deserving of trust.

REFERENCES

- [1] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N. Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, and Eric Horvitz. 2019. Guidelines for human-AI interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, Paper 3, 1–13. DOI:<https://doi.org/10.1145/3290605.3300233>
- [2] Lisanne Bainbridge. 1983. Ironies of automation. *Automatica* 19 (1983), 775-783.
- [3] Alan Dix. 2002. Beyond intention-pushing boundaries with incidental interaction. In *Proceedings of Building Bridges: Interdisciplinary Context-Sensitive Computing*, Glasgow University, Vol. 9.
- [4] Peter Fröhlich, Matthias Baldauf, Thomas Meneweger, Manfred Tscheligi, Boris de Ruyter, and Fabio Paternó. 2020. Everyday automation experience: a research agenda. *Personal and Ubiquitous Computing* 24 (2020), 725-734. <https://doi.org/10.1007/s00779-020-01450-y>
- [5] Roberto Hoyle, Srijita Das, Apu Kapadia, Adam J. Lee, and Kami Vaniea. 2017. Was my message read?: privacy and signaling on Facebook Messenger. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. Association for Computing Machinery, New York, NY, USA, 3838-3842.
- [6] Christian P. Janssen, Stella F. Donker, Duncan P. Brumby, and Andrew L. Kun. 2019. History and future of human-automation interaction. *International Journal of Human-Computer Studies*, 131, 99-107.
- [7] Wendy Ju and Larry Leifer. 2008. The Design of Implicit Interactions: Making interactive systems less obnoxious. *Design Issues* 24, 3 (2008), 72-84. arXiv:<http://dx.doi.org/10.1162/desi.2008.24.3.72>
- [8] Microsoft Viva Topics. Retrieved February 27, 2021, from <https://www.microsoft.com/en-us/microsoft-viva/topics/overview>
- [9] Wanda Orlikowski. 2002. Knowing in practice: enacting a collective capability in distributed organizing, *Organization Science* 13, 3 (2002), 249–273.
- [10] Barış Serim and Giulio Jacucci. 2019. Explicating "implicit interaction": an examination of the concept and challenges for research. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, Paper 417, 1-16. DOI:<https://doi.org/10.1145/3290605.3300647>
- [11] Suvi Silfverberg, Lassi A. Liikkanen, and Airi Lampinen. 2011. "I'll press play, but I won't listen": profile work in a music-focused social network service. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work (CSCW '11)*. Association for Computing Machinery, New York, NY, USA, 207-216.
- [12] Rayoung Yang and Mark W. Newman. 2013. Learning from a learning thermostat: lessons for intelligent systems for the home. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '13)*. Association for Computing Machinery, New York, NY, USA, 93-102.