

Ways of Seeing and Being Seen: People in the Algorithmic Knowledge Base

Ida Larsen-Ledet & Siân Lindley, Microsoft Research Cambridge

Introduction

Ever since computers enabled large scale archiving and structuring of an organization's knowledge to make it accessible and (re-)usable for employees, digital knowledge management has been both a promise and a challenge for organisations. With a keen eye set on making better and more efficient use of knowledge, and reducing duplication and redundant work, organisations have long struggled with making their overwhelming amounts of dispersed information useful. But this task itself is overwhelming precisely because of the volume and dispersal of content [1]. These days, cloud solutions such as Document360 and Confluence are in widespread use, and between communication and collaboration tools like Workplace and Slack, and productivity tools like Office 365, information is not only created by workers and saved across a vast ecology of artifacts; it is also automatically generated from people's activities. The voluminous and continuously changing and expanding nature of organisational knowledge makes knowledge management an obvious application area for artificial intelligence, which can offload the manual labour that usually goes into building and curating knowledge bases [2]. We are seeing the use of machine learning for compiling, summarizing, and surfacing knowledge within an organisation in commercial products like Viva Topics [3].

These kinds of algorithmic systems enable knowledge management to move beyond putting things in the right folders to the (re-)shaping and surfacing of knowledge, not only from published content but also from the ongoing work that people do. With this, a knowledge base is no longer just a repository but becomes an active part of people's work [2]. When working to understand and develop support for knowledge management, we are thus faced with questions not only about how to categorize and present knowledge but also about how knowledge *happens* [4, 5, 6]. In Orlikowski's words, organisational knowledge is most usefully viewed as a process, i.e., with a focus not on knowledge artifacts but on the knowing that emerges in "the ongoing and situated actions of organizational members" [4]. These actions include interactions between people and the contextual understanding used when navigating those interactions, such as knowing who to ask about what and how to ask them [6].

In this position paper, we discuss what it will take for knowledge bases to scaffold processes of knowing [4] and expand beyond systems for sharing knowledge items. We propose that interaction patterns [7] may be used to reframe the focus of knowledge base design onto the (inter-)actions that enable active knowing. We take *knowing who to ask* as a case in point, discussing the challenges of associating specific individuals with knowledge or content.

Background: Interviews with Knowledge Workers

Our thinking takes outset in a recent interview study of how knowledge workers characterise the knowledge they work with. The participants were knowledge workers (people whose main capital is the application and creation of knowledge) and knowledge managers (people who are responsible for managing an organisation's knowledge base). In interviews with 25 knowledge workers (10 women, 15 men) and 5 knowledge managers (2 women, 3 men), we explored what attributes of the knowledge are essential for their organisation and the work that happens in it. We explored this through a mapping exercise, in which the participant selected different types of knowledge that were relevant to their work – such as "contract", "event", and "client". In one of the ensuing steps of the exercise, the participant would provide concrete examples of different types of knowledge and describe what attributes they associated with these; which of these attributes were most essential; and what purpose they would serve. The value of *knowing who to go to* or *who to ask*, which we focus on in this paper, emerged as a prominent topic from the exercise. While the paper originates in this empirical work and addresses one of our primary findings from it, the reflections presented here are largely theoretical.

Knowing Who to Ask

Clyde¹ works as an engineering manager. Because of the rapid development of their product, their data model is constantly evolving and is, hence, not well-documented. Clyde therefore often needs to identify and reach out to "the right owner" to understand how a given part of the product has been modelled.

Eric is a project manager for a utility company. Part of his job is to monitor the status of electrical equipment, but their data warehouse doesn't always contain the data he needs. In these cases, Eric will reach out to "the relevant team" and ask if more data is available and if he can get it in an Excel file.

Much of the knowledge (and knowing) of an organisation is connected to people who work there [4, 8, 9]. This knowledge is not only contained in tangible content but also in practices and knowhow [4, 8, 5, 10] and, being particular to different communities of practice in the organisation, it becomes actionable through knowing "what to attend to and what to avoid" [8]. Making assessments like this when putting knowledge to work requires situated and occasioned understanding of circumstances and local knowledge [6]. Brown and Duguid [8] point to the difficulty of transferring this kind of understanding as a reason why putting knowledge to use in new contexts is often challenging for organisations. Content may be transferred from one work setting to another, but that does not necessarily mean that actionable knowledge, as it is tied to the practices of the work setting, travels with it [4]. A knowledge base, even if powered by artificial intelligence, does not understand context in the same capacity as a person and so cannot "broker" or "translate" that knowledge to the same degree as a conversation between people [8]. One person asking another can produce an interaction in which actors knowingly ask and answer – what Randall et al. [6] call *recipient-designed* recalling. Lists of "related people", or other ways of helping workers find people *who know*, are thus a way to enrich knowledge bases: Not only for workers like Clyde and Eric to have content explained by the person who created it or being able to ask for knowledge in a certain format but also to resolve confusion or probe deeper to learn more. Moreover, knowing what people work with the knowledge in question may provide useful cues for interpreting and understanding the knowledge. Sometimes, work will also rely on knowledge that cannot be documented in the knowledge base, such as the informal classifications of unemployed applied by caseworkers in a Danish job centre [11]. That knowledge must be passed on directly between people. These examples support Brown and Duguid's [8] case that people are needed for translating and brokering knowledge. Identifying different kinds of *knowledge actions*, Lindley et al. [12] likewise suggest that a knowledge base should help workers find both information and people who can provide information, and that the knowledge base therefore needs to contain information about *who knows what*.

How to Know Who to Ask

Helping people know who to ask may be an initial step in facilitating knowledge sharing between colleagues [13]. There are many layers to be added onto this basic step, however, in terms of the way asking and answering takes places. For example, asking may often not be a one-step process: If you don't "just know" (as some of our participants put it) who to ask, you might reach out to someone who can point you to the right person. This not only enables you to find *a* person to ask; since the answers are provided in the specific context of what and why you need the person, talking to a middleperson might get you additional knowledge that helps you know when and how to ask, and what you might expect in terms of them getting back to you. Furthermore, the middleperson might know about the other person's situation as well, being able to direct you to someone who is, e.g., less likely to be busy (related to what Orlikowski refers to as "knowing of each other" and each other's commitments [4]).

¹ The names used for participants are pseudonyms.

A list of related people serves the same purpose as the middleperson. But listing related people in a knowledge base like dishes in a restaurant menu is not contextual in the same way and does not come with contextually relevant additional information, such as how/whether it is appropriate to approach each person about a given subject. This relates to classic discussions about the way information technology influences practices and the way work is carried out, and how technology can be disruptive by, directly or indirectly, changing the way people engage with each other in the workplace (see, e.g., [14]). So how can a knowledge base more actively mediate knowledge exchanges instead of merely being a directory? A small improvement may be to add information such as roles, working hours, or current status (like the busy/away badges we know from chat applications like Slack and Teams). But in an automated knowledge base, the first challenge is deciding who to list in the first place. Modern knowledge bases may have access not only to content but also to information about who worked on that content, what they did with it, and when, as may be the case for systems like the ones mentioned in the beginning of this paper where knowledge is collected not only from repositories but also from workers' activities. It appears obvious to use this kind of information to infer who might be the right person to ask about a given topic [2]. However, this will inevitably involve implicit assumptions about what "right" means (and what it does not mean). As implied above, taking "the right person" to simply mean, e.g., "someone who has something to do with this topic" might be inaccurate or inappropriate for multiple reasons; from that person not having the time, to them not understanding enough about the work that the person asking needs the knowledge for (see also [13]). This is to say that "right" is contextual and can change from one moment to the next. On the one hand, if any technology can deliver the necessary adaptability and capacity for sifting through the organisation's knowledge, it will be artificial intelligence. On the other hand, exploring how to replace or supplement a human middleperson shows us that defining the right problem and the right assumptions is not straightforward.

Just as importantly, we may consider the consequences for the person being listed. Does being listed as a related person in a knowledge base entry about a subject imply that you're happy to answer questions from work acquaintances, or perhaps even strangers, about that subject? Wilkins et al. [2] point out the risk of disempowering workers by treating them as a resource to be mined for knowledge: Workers may experience having roles and obligations thrust upon them that, while not enforced by the system, have roots in the way they are represented by the system and how co-workers interpret and act on that. Listing a person as being related and, i.e., someone to ask does not produce the knowing of each other achieved through personal connection and direct interaction [4]. When people are represented as resources, especially if the representation lacks contextual depth as to why, how will others interpret and decide what actions are appropriate to take towards those people? The adaptability of algorithmic systems does not inherently solve this but must be utilized in the right way. Given the challenge (or impossibility) of determining who the right people to suggest are, we may look to Orlikowski and redirect our focus away from the knowledge (what and who) and over to the knowing (how) [4].

Seeing Knowing as (Inter-)Actions

As we have discussed above, organisational memory, like other kinds of remembering, is socially organised [6]. But where does social organisation fit into current organisational knowledge bases? Traditional storage models decontextualize knowledge [6] because they focus only on the knowledge and neglect the people and interactions that activate it [4, 5]. Exploring how different metaphors for a library influenced the design of library computer support [15], Halskov Madsen describes how seeing the library as a warehouse *foregrounded* the books and *backgrounded* the borrowers, resulting in computer systems that prioritized stock keeping, accounting, and book retrieval, despite their original promise being improved service for borrowers [16]. As an alternative, Halskov Madsen outlines a meeting place metaphor, in which borrowers are foregrounded and where librarians are reconceived as consultants or mediators. Halskov Madsen's explorations demonstrate the power of *seeing as* [7, 16, 15], and how identifying and challenging current framings can be conducive to innovation.

Listing people like any other resource in the knowledge base makes knowing who to ask into a question of using a person like one would any other resource. Reframing the knowledge base in terms of socially scaffolded knowing, or recipient design, would make it into a question of people, plural, actively working out ways of asking and answering. Here, we need to keep in mind that using artificial intelligence does not mean that we have to take away people's room for acting intelligently themselves. As Ju and Leifer [7] note, technological intelligence needs to be supplemented by *sociable design* for technology to function in uncertain situations (which a lot of situated work arguably is). This requires us to understand how people accomplish knowledge sharing in practice [4, 6]. That is, we should turn our attention to action rather than "infrastructure, objects, skills, or dispositions" [4]. If we, in addition to asking what kinds of knowledge people look for, ask: "What steps do people take to find knowledge, and in what ways do they use its different attributes?" we can understand how they carry out the skilled activity of obtaining and using the knowledge they need, and how we can help them do that. To make this manageable in a design setting, it may be useful to capture those practices in *interaction patterns*: Ju and Leifer [7] propose mapping technological capabilities against situated human interaction, to envision ways of guiding situated interaction with technology. These interaction patterns can be used as lenses for seeing a design problem in terms of interactions between people [7], e.g., seeing a knowledge base in terms of ways of knowing. Mapping to concrete patterns of interaction lets designers determine the roles, settings, and sequences of action [7] that enable the successful accomplishment of work. Among other things, explicitly identifying what roles are at play in socially accomplished knowing could help recast related people from secondary participants in an information seeker's quest to collaborators whose working practices and circumstances are relevant. Like metaphors, the labels we apply to people can foreground or background different features of a situation and the roles at play [9, 16].

Conclusion

While broadly applicable, these suggestions are especially relevant for systems where it is not humanly feasible to curate each bit of content to help it be put to work in every conceivable (and inconceivable) situation. Our concerns reach beyond artificial intelligence – they have to do with computer representations in general – but are especially pronounced in algorithmic systems because of how much is left up to the machine. The issue may even be exacerbated by the term "intelligence": The system is intelligent and so must have acted intelligently when collecting and categorizing the information being presented. Our expectations of the system's capabilities inevitably shape the way we act on what the system provides, regardless of the degree to which we are conscious of those expectations. Algorithmic systems like Viva Topics, in which knowledge managers and workers can curate the information compiled by the system [3], may be an interesting case study. For instance, workers may find themselves in an awkward position of not knowing whether the list of related people has been curated by someone who knows, or has the authority to declare, that those people are the right ones to contact (whatever the understanding of "right"), or whether the list is based on the system's assumptions and may not have been filtered for local practices. As the conundrums around the notion of "the right person" exemplify, even automated knowledge bases are not choice-less, and the choices have consequences. In the case of related people, they have consequences for the way people are seen *and not seen* [9]. How might being included in or excluded from the list of related people influence workers' opportunities to have impact [12]? How will workers change the ways they perform their work to respond to these consequences of algorithmic categorisation [9, 12, 17]? How will their concerns about being seen in the right way affect measures of success? These questions highlight the reciprocity between technology and practice. Of course, it is not by definition problematic if a "restaurant menu" of related people encourages workers to reach out directly to strangers, but to paraphrase ethics and privacy scholar Helen Nissenbaum: That technology can change what is considered acceptable should not by default be seen as a bad thing, but as an outset, entrenched norms and existing practices should be respected [18]. Algorithmic systems may produce high measurable accuracy or coverage, but user research and UX design have vital roles to play in making knowledge actionable and making sure it is actionable *in the right ways*.

References

- [1] H. Tsoukas, "The Firm as a Distributed Knowledge System: a Constructionist Approach," *Strategic Management Journal*, vol. 17, no. Winter Special Issue, pp. 11-25, 1996.
- [2] D. Wilkins, S. Lindley, M. Meijer, R. Banks and B. Burlin, "Designing Ai Systems That Make Organizational Knowledge Actionable," *Interactions*, vol. XVII, no. 6, November - December 2020.
- [3] Microsoft, "Microsoft Viva Topics," Microsoft, [Online]. Available: <https://www.microsoft.com/en-us/microsoft-viva/topics>. [Accessed 25 April 2022].
- [4] W. Orlikowski, "Knowing in Practice: Enacting a Collective Capability in Distributed Organizing," *Organization Science*, vol. 13, no. 3, 2002.
- [5] D. Randall, J. Hughes, J. O'Brien, M. Rouncefield and P. Tolmie, "Memories are made of this: Explicating organisational knowledge and memory," *European Journal of Information Systems*, vol. 10, pp. 113-121, 2001.
- [6] D. Randall, J. O'Brien, M. Rouncefield and J. A. Hughes, "Organisational memory and CSCW: supporting the 'Mavis Phenomenon'," in *Proceedings Sixth Australian Conference on Computer-Human Interaction*, 1996.
- [7] W. Ju and L. Larry, "The Design of Implicit Interactions: Making Interactive Systems Less Obnoxious," *Design Issues*, vol. 24, no. 3, pp. 72-84, 2008.
- [8] J. S. Brown and P. Duguid, "Organizing Knowledge," *California Management Review*, vol. 40, no. 3, pp. 90-111, 1998.
- [9] G. C. Bowker and S. Leigh Star, *Sorting Things Out: Classification and Its Consequences*, MIT Press, 2000.
- [10] D. J. Wilkins, S. H. Muralidhar, M. Meijer, L. Lascău and S. Lindley, "Gigified Knowledge Work: Understanding Knowledge Gaps when Knowledge Work and On-Demand Work Intersect," *Proceedings of ACM Human-Computer Interaction*, vol. 6, no. CSCW1, 2022.
- [11] A. C. M. Petersen, L. R. Christensen, R. Harper and T. Hildebrandt, "'We Would Never Write That Down': Classifications of Unemployed and Data Challenges for AI," *Proceedings of the ACM on Human-Computer Interaction*, vol. 5, no. CSCW1, 2021.
- [12] S. Lindley, D. Wilkins and B. Burlin, "Actions and their Consequences? Implicit Interactions with Workplace Knowledge Bases," in *Automation Experience at the Workplace (workshop in conjunction with CHI'21)*, 2021.
- [13] M. S. Ackerman, "Augmenting the Organizational Memory: A Field Study of Answer Garden," in *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*, 1994.
- [14] J. Grudin, "Groupware and Social Dynamics: Eight Challenges for Developers," *Communications of the ACM*, vol. 37, no. 1, pp. 92-105, 1994.

- [15] K. Halskov Madsen, "Breakthrough by Breakdown: Metaphors and Structured Domains," in *IFIP WG 8.2 conference: Information Systems Development for Human Progress in Organizations*, Atlanta, 1987.
- [16] K. Halskov Madsen, "A Guide to Metaphorical Design," *Communications of the ACM*, vol. 37, no. 12, pp. 57-62, 1994.
- [17] M. S. Ackerman, "Definitional and Contextual Issues in Organizational and Group Memories," in *Proceedings of the Twenty-Seventh Hawaii International Conference on System Sciences*, Wailea, Hawaii, USA, 1994.
- [18] H. Nissenbaum, "Privacy as Contextual Integrity," *Washington Law Review*, vol. 79, no. 1, pp. 119-158, 2004.