

Title	Generative AI at work: mundane and productive articulation
Authors	Larsen-Ledet, Ida
Publication date	2026
Original Citation	Larsen-Ledet, I 2026, 'Generative AI at work: mundane and productive articulation', Paper presented at 2026 CHI Conference on Human Factors in Computing Systems, Barcelona, Spain, 13/04/26 - 17/04/26 pp. 1-6. https://doi.org/10.1145/3772363.3798602 ;conference
Type of publication	Paper (Peer reviewed)
Link to publisher's version	10.1145/3772363.3798602
Rights	cc_by
Download date	2026-08-28 06:17:52
Item downloaded from	https://hdl.handle.net/10468/18688

Generative AI at Work: Mundane and Productive Articulation

Ida Larsen-Ledet

ilarsen-ledet@ucc.ie

School of Applied Psychology, University College Cork
Cork, Ireland

Abstract

This paper proposes a distinction between productive and mundane articulation in people's interactions with generative AI. While generative AI produces useful results for certain tasks and domains, for many types of work, experiences with AI are still dominated by generic and superficially relevant responses. People are often required to provide context and clarifications through several iterations to get usable output. This paper argues that generative AI will not become a true asset in the workplace unless it allows people to focus on *productive articulation* that helps them reflect on and refine the objectives of their work, by reducing the need for *mundane articulation* that makes up the rote work of setting up fundamental context for generative AI to deliver useful results.

CCS Concepts

• **Human-centered computing** → **Natural language interfaces**; *HCI theory, concepts and models*; *Interaction design theory, concepts and paradigms*.

Keywords

generative AI, human-AI interaction, AI articulation, relevance, personalization, context, knowledge work

ACM Reference Format:

Ida Larsen-Ledet. 2026. Generative AI at Work: Mundane and Productive Articulation. In *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3772363.3798602>

1 Generative AI at Work

Generative AI (genAI) seems to be the new medium for knowledge consumption and use. With generative AI, tools that draw information from multiple sources and combine it on the spot to present to a user are becoming commonplace [43]. So far, however, the user experiences they offer remain fairly rudimentary. They are good, or at least reasonable, at providing well-scoped responses to well-specified questions or prompts. But the more subjective or specialized the request, the more evident the generic nature of genAI becomes. At first glance, the value of genAI comes from its ability to summarize, to generate, and to fill in gaps¹. But the

real value and, in particular, the novelty of genAI comes from it being responsive – that is, generating things that are *suitable* for a purpose or a situation [43, ch. 4]. Without the summarization, generation, and gap-filling being relevant, they would not be as useful, at least not for most work activities.

1.1 AI as Part of Knowledge Infrastructures

Recent years have seen the uptake of organizational knowledge bases that are automatically constructed and managed with the aid of machine learning and AI [e.g., 18, 19, 45], and the introduction of generative AI to workplace software is providing new ways of utilizing these knowledge infrastructures. Both of these developments have enabled leaps from passive knowledge stores to active surfacing of knowledge [44], bringing us closer to the goal of “connect[ing] knowledge workers with the right set of knowledge resources or people, at the right time” (O'Dell and Davenport as quoted by Jarrahi et al.) [18, p. 97]. GenAI, in particular, is changing the use and reuse of knowledge in organizations by changing the way it is surfaced [43]. The vision seems to be that genAI will be integrated in all manner of workplace applications, or at least office applications ranging from data analysis to word processing and meeting software. With this, knowledge that previously sat much more passively in knowledge bases, mainly to be consumed, may now more actively influence people's work as it happens.

By (organizational) *knowledge infrastructure*, I refer to knowledge storage, management, sharing, and (re)use of knowledge through knowledge bases, information retrieval, and expertise finding. I use the word *infrastructure* to denote the interplay between multiple tools, technologies, and practices [23, 33] and how they enable knowing (the reuse and application of knowledge) [1, 3, 28]. I use the term *knowledge infrastructure* rather than *organizational memory* [4] to describe the combined technological, practical, and social means for organizational remembering because technology in itself is not memory but an *infrastructural enablement of memory* as something active and constructive [5]. That is, I am not talking just about (passive or active) knowledge bases but about technology that facilitates the human acting involved in remembering [4] and using knowledge. This is what good genAI should do – not replace human memory, interpretation, or knowledge use but facilitate them as action.

Information in an organization often has multiple potential use scenarios, and there is lots of it. Employees constantly assess what information to pay attention to, what to make active use of, and what to disregard – in other words, what is relevant to them and to what they are currently doing. As genAI selectively synthesizes information for people, how can those people be confident that they are actually getting to (all of) the knowledge they need and want for what they are doing? Even beyond information finding specifically, genAI's connection with and place(s) within knowledge

¹Filling in gaps in a broad sense of the term: For example, as the Copilot in Microsoft's Visual Studio Code completes method bodies based on a user's instructions or code comments, it is essentially filling in gaps laid out by the user.



This work is licensed under a Creative Commons Attribution 4.0 International License. *CHI EA '26, Barcelona, Spain*

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2281-3/2026/04

<https://doi.org/10.1145/3772363.3798602>

infrastructures are important, as all of genAI’s capabilities are tied in with information through regurgitation and remixing of it. This paper discusses what it means to situate AI in people’s work, as genAI applications promise to weave knowledge surfacing and remixing into the fabric of knowledge work.

2 Background

2.1 The Work to Make AI Work

People take to generative AI and its ability to select and summarize as a way to reduce the overhead of finding or consuming information [35, 47]. But it often misses the mark because of mistaken or lacking assumptions [e.g. 22]. This shows that being over-confident in assumptions about the user leads to poor or less-than-useful user experiences.

Although a frequent assumption is that genAI enables people to spend more time on “core” work activities and less time on articulation work and tedious rote work [e.g. 17, 35, 46] (which, arguably, might depend on how closely your work lends itself to what current genAIs do well), being productive with genAI requires people to skillfully articulate themselves in the right way, for instance specifying details that would be implicitly understood by a coworker [35, 36] (e.g., that a slide deck should follow the institutional template), and proposed solutions have so far hinged on people doing this explicitly [e.g. 20, 21].

2.2 Relevance

The notion of relevance is essential to organizational knowledge infrastructures, as information only becomes usable knowledge when that information is meaningful to a person and a situation [8]. When viewing organizational knowing as emergent from the situated skillful acting of people [28], the connection to person and situation also depends on what that person is doing with the information. As such, information’s status as knowledge depends on how relevant the information is to the person and their activity. In the fields most often interested in studying information relevance, such as information retrieval, relevance was originally conceived as topicality: in essence, the match between terms in a search query and terms in a search result [30]. It has long since become clear that relevance is *contextual* [29], and that not all factors pertinent to relevance can be codified [29, 30]. With this in mind, when we consider genAI as machine that relies on textual semantics, the experience of surface-level relevance becomes less surprising. Its textual operations may be quite advanced, but the effects of primarily topical relevance are still felt [21, 35].

2.3 Personalization and Context

Personalization for information consumption and knowledge use can consist both in adapting the *information* that is provided to people and in adapting the way that information is *presented* and the available modes for *interacting* with it [9, 27]. This adaptation will be based on information about the person and/or their situation [e.g., 9, 14]. Underlying many such proposals is the premise that what is explicit – often a search query – is not enough for determining what a person needs (see also contextual relevance above) [40]. Teevan et al. [39], for example, demonstrate that the same query can carry different intent depending on who utters it. This underlines the need

to consider not just the person in isolation but their broader context. Bannon and Kuutti [4] argue that the remembering and reuse of information is crucially dependent on context, and AI systems underscore context as a central notion: Recent work has highlighted the problems that result when AI systems are unaware, or only partially aware, of the context in which outputs are surfaced or produced [13, 37], and other studies have shown that taking broader contextual information into account can improve AI’s ability to predict people’s information needs [41, 42].

3 Situating AI in People’s Work

So, context and clarification help genAI produce better results. But so far, attempts to build clarification into genAI applications through explicit input fail to distinguish between things that one would expect to need to specify and things that should be implicitly understood by an “intelligent” agent. For genAI’s output to feel not only surface-level relevant but also personally and situationally relevant, we need to distinguish between articulation that is a productive part of work and articulation that is simply mundane.

This is not just about who the person is or what information they need but also what needs to *happen*. Usefulness depends on how fit for purpose the possibilities for action are. Current modes of (inter)action in most genAI tools are relatively basic, as are the places they appear in: Most of them are restricted to conversational interfaces [43, ch. 4] or, at best, confined to making things happen in a single app. This does not match how people work [6], and it effectively *decontextualizes* use by siloing the AI. Interaction with genAI should be contextualized in relation to people, activities, and collaborative settings, not just in relation to whatever app is currently in use.

3.1 Mundane and Productive Articulation

To deliver the most useful response, the system or tool should not only understand the explicit content of a request or statement but also the intent behind it [40, 43]. A request such as: “who should I ask about the Planning department’s procedure for new build surveys?” should produce different responses depending on who is asking and what situation they find themselves in. However, many current genAI tools seem to operate with a superficial notion of context that relies on explicit information such as the user’s organizational or physical location, or what software application the given instantiation of genAI “lives” in (or simply the preceding utterances in a conversational interaction [43, ch. 3]).

If genAI systems do not become capable of genuine contextualization, people will continue to be forced to articulate the most mundane aspects of their situation before getting to anything useful. Many readers have probably tried the back-and-forth it takes to get genAI tools to carry out a request in a way that is right for them. A lot of this involves getting the tool to “understand” you and any relevant circumstances [36]. We can draw a parallel to the articulation work that is required in collaborative work: For multiple actors to work together effectively, they need to establish and maintain common ground, and to understand responsibilities and division of work, etc. [31]. This articulation work is not by definition good or bad; it is simply a part of collaborative work. Similarly, genAI requires some articulation of intent and context

on the user’s part, which requires the person to figure out both *what* to tell the system and *how* to tell it. This articulation *can* be productive, just as it can be productive when we explain a problem to a coworker (or a rubber duck on our desk) only to come up with an idea for a solution ourselves without the coworker needing to speak a word – but articulating mundane aspects (such as the fact that if I need ideas for how to structure an article, I mean an academic article and not a news article), especially repeatedly, is usually not. If genAI requires people to spend their energy on *mundane articulation*, it is likely to become an additional cognitive burden rather than fulfilling the promise of freeing up people’s time to let them focus on what matters to them.

If, instead, genAI’s responses and suggestions are contextually anchored, the demand on people to articulate contextual aspects will be reduced, allowing them to turn their focus to *productive articulation* that helps them reflect on and develop their conceptualization of the task they want to do or the problem they are trying to tackle. *In sum: (Appropriately) Anchoring people’s use of AI in their context (person and situation) can allow them to focus on productive articulation rather than mundane articulation.*

4 So, what is the mundane?

... or: what is it mundane to have to articulate (over and over)? To paraphrase the above, the mundanely articulated is that which adds no value beyond ensuring that outputs are suitable for the person and the situation, while the productively articulated is that which progresses not only the generation of output but also how the work is conceptualized. A parallel can be drawn to Shah and Bender’s notion of “beneficial friction”, which acknowledges that information access is not just about getting an answer but also about sense-making [32]. My purpose here is not provide a concrete list of what is mundane and what is productive, partly as the answers will likely be contextually dependent themselves. My purpose is to point to the crucial distinction between productive and mundane articulation, begin a dialogue, and propose a direction of inquiry.

In essence, I propose mundane and productive articulation as lenses to examine (gen)AI systems with, for purposes of both understanding and design. A possible line of inquiry in this regard is whether it is possible at a general level to map different aspects of a situation to a mundane/productive articulation distinction or axis, and if (likely) not, how dependent such a mapping is on the setting, the domain, the application area, and so on. It may (or may not) be that it is always less mundane to have to state what you are working on right now than to reiterate that you are a researcher and not a marketing director or a student. Another part of such an inquiry is to decide if mundane and productive articulation are most usefully conceived of as discrete categories or as ends of a spectrum. This, too, may prove to be dependent on other factors. Figure 1 summarizes this mapping challenge.

To have some scaffolding for this inquiry, it would be helpful to have a properly nuanced understanding of context. Context is a contentious notion in HCI, and one definition or model of it may be suitable for one use case and inappropriate for another, making it impossible to definitively say what “context” is comprised of [7, 12, 16]. Abstracting away from an enumeration of concrete components of context, with the intention of focusing on the *purpose* of context

information in this inquiry, one trap we might fall into is to assume that some aspects of context are completely situational – such as who else is in the room or what they are trying to achieve together – and others are a parcel of information that always travels with the person and is always present, such as a person’s job or educational background. It would be easy, then, to assume that it is mundane to tell the system what your job is and productive to tell it what you are trying to achieve. But these things may not actually be so straightforward to separate. For example: Articulating what or how much you already know about a subject, to avoid the AI over-explaining, can be both difficult and tedious – on the other hand, it can (as mentioned) be very productive to voice your thoughts and assumptions out loud, including what you already know about a problem and what knowledge you might apply to solve it. To tackle this, we should factor in that circumstances and characteristics can be more or less defining for a situation.

Even relatively constant or slow-evolving characteristics come into play in a situational manner. For instance, the knowledge a person draws on in a given situation will often be different from the knowledge they draw on in another, even though the knowledge they possess as a whole obviously does not completely change from situation to situation – a person might be employed as a quality assurance inspector, drawing on specialist knowledge of biochemistry, while also chairing an employee welfare committee, drawing on broad insight into organizational culture and procedures. A prompt like: “draft an email summarizing these complaints about test equipment” will carry significant underlying implications depending on whether the email is going to the whole organization, just the quality assurance team, or just management, and depending on what capacity the person is acting in; specialist impacted by test equipment issues or leading representative for employee welfare. With this framing, it is clear that all aspects of context are situationally entwined and that it is not trivial to specify (certainly not deterministically) which aspects of a person’s context are mundane to articulate and which are productive.

Another question we may ask is whether this is only about what users should be able to leave implicit; or if what is mundane and what is productive for the user to articulate should also make a difference for how that information influences the AI: When a genAI is used to compose a report for a project manager, should mundane facts, such as their product manager title and that they work in the security division, be used or treated differently by the AI than more situational, and potentially more productively articulated, factors such as an underlying strategic goal of getting a related project extended? Might the more mundane, for example, best be used in a dynamic system prompt, while the less mundane is use to enhance automatic prompt reformulation² (or the other way around)?

A whole other question is how much information people are comfortable with a given system “knowing” about them.

4.1 Exploring the Distinction

As outlined, determining which circumstances and characteristics are mundane to articulate and which are productive, or the extent to which they are one or the other in a given situation or setting, is

²The process of rewriting a prompt or query to improve expected effectiveness before running it [49].




Figure 1: The articulation of intent and context can be productive, but articulating the mundane is not. For AI to actually free up people’s time, this distinction is necessary. But the mapping (illustrated here using dashed arrows) of circumstances and characteristics onto mundane and productive articulation, if at all possible, is an open question (indicated by overlaid question marks); as is whether mundane and productive are discrete categories or ends of a spectrum (illustrated through separated boxes and a gradient-filled arrow, respectively). I encourage this as a line of inquiry in future work on human-genAI interaction.

not straightforward. An obvious avenue is to study people’s articulation in genAI interactions, using approaches such as think-aloud or diary studies to elicit thoughts about the things they find necessary to specify to the AI, and employing mundane and productive articulation as a framing when analyzing participants’ experiences. Even without targeted studies, it would be instructive to apply the productive/mundane framing as a lens onto existing literature that identifies challenges [e.g. 22, 35] and compare across cases.

Considering mundane and productive articulation may also serve a generative purpose in the design of genAI tools or features, such as prompt assistance or task- or domain-specific genAI tools; for example by guiding explorations of what assumptions should or should not be made by the AI, or what clarification to seek from the user and how to integrate clarification into interactions in a productive way.

5 GenAI is an Answer, Not the Answer

That I discuss how to make genAI better should not be taken as an argument for genAI to be used everywhere or all the time. GenAI now seems to be part of the future of work and, as such, should be considered in research on work, particularly given its links with organizational knowledge use (Microsoft’s M365 Copilot, for example, uses an organization’s files to answer questions and surface potentially useful files to the user [26]). However, it is both worthwhile

and important to consider that the way to create more relevant knowledge experiences at work can be *better design*, with or without (better) AI.

5.1 Is Generative AI Worth It?

Any suggestions for improving AI also become suggestions for advancing AI’s role in the world, and I am ambivalent about contributing to this. If the choice is between better AI and what we currently have, better is the obvious choice. But Sturdee et al. rightfully argue that “the potential negative consequences of research should be questioned, critiqued, and discussed as part of the publication and peer review process” [34, p. 1]. (Generative) AI comes with a range of potential and certain costs. There are consequences for people using it or working in an organization where it is used. There are also consequences for other people, both before that use and as a consequence of it. Consequences *before* use arise (and have already arisen), for instance, during model training and the creation of datasets for that training, e.g., when data labellers are underpaid yet have little choice but to do the work [15, 38], in many cases without being part of the population that reaps the largest benefits [48]. Consequences *of* use could be when people’s lives are influenced directly or indirectly by decisions and actions taken in the course of work mediated by AI-powered tools (a very concrete parallel is the use of predictive algorithms in, for example, social

services and police work [25]). It could also, more indirectly, be the climate harms and displacement [2] caused by the environmental impact of generative AI. Environmental consequences stem, for instance, from the high amounts of energy used when creating and training models [10]. Importantly, this cost is not one-off, as (like in any design and development process) getting to the right functionality and performance will require multiple rounds of iteration. And even if existing models can be (re-)used to implement the principles, running the models also costs both energy [24] and water for data center cooling [10], which would quickly mount up with the high numbers of users of knowledge work software (whether they use it for work or personal uses): A comparative study of the cost of inference using different generative models found that the cost of prompting outweighs the cost of training and fine-tuning within weeks or months, particularly as more and more people begin to use generative AI [24]. While it may be possible to choose a model with a lower cost associated with its development and/or its use, there is (unsurprisingly) a trade-off between accuracy and how environmentally (un-)sustainable different models are [11]. It is important to assess whether a poorer implementation using genAI could be worse than no implementation at all, for example, as people would need to iterate more to get what they need, potentially resulting in costs the same as or higher than with a better model. But to flip the coin yet again: Would a more effective system mean briefer use, or would people merely fill in the “freed-up” time with more use of the system (similar to when the energy efficiencies created for bitcoin coincided with greater demand and mining [10])?

This – yet another paper about how we can make genAI better – should not be read as an uncritical endorsement of (generative) AI. It is definitely worthwhile to consider what difference it could make to use (generative) AI or not to use it; what model(s) to use if AI is used; and whether the same things can be achieved sufficiently without it. Importantly, this is not just about whether the use of generative AI (or computing in general) does damage but about whether that damage is excessive compared to the social, economic, and environmental benefits [10]. So, not only is it necessary to assess whether the immediate goals of implementing this, or any other, paper’s ideas would be worth it but also what the benefits would be in a grander scheme – do they really measure up against the potential negative impact? The ideas in this paper must not be used uncritically. This should go without saying for any paper making recommendations about technology, but I am not sure it always does.

Acknowledgments

Thank you to Sean Rintel and Lev Tankelevitch for showing an interest in the conceptual distinction between mundane and productive articulation when I used it in a presentation in 2024. This encouraged me to develop my thinking on it further.

This work was partially funded by the European Union under Marie Skłodowska-Curie Actions Postdoctoral Fellowship grant agreement no. 101110480. Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or the European Research Executive Agency (REA). Neither the European Union nor the granting authority can be held responsible for them.

References

- [1] Mark S. Ackerman, Juri Dachtera, Volkmar Pipek, and Volker Wulf. 2013. Sharing Knowledge and Expertise: The CSCW View of Knowledge Management. *Comput. Support. Coop. Work* 22 (2013), 531–573. doi:10.1007/s10606-013-9192-8
- [2] Reazul Ahsan, Jon Kellett, and Sadasivam Karuppannan. 2014. Climate Induced Migration: Lessons from Bangladesh. *The International Journal of Climate Change: Impacts and Responses* 5, 2 (2014), 1–15. doi:10.18848/1835-7156/cgp/v05i02/37204
- [3] Liam Bannon and Susanne Bødker. 1997. *Constructing Common Information Spaces*. Springer Netherlands, Dordrecht, 81–96. doi:10.1007/978-94-015-7372-6_6
- [4] L.J. Bannon and K. Kuutti. 1996. Shifting perspectives on organizational memory: from storage to active remembering. In *Proceedings of HICSS-29: 29th Hawaii International Conference on System Sciences*, Vol. 3. 156–167 vol.3. doi:10.1109/HICSS.1996.493187
- [5] Liam J. Bannon. 2006. Forgetting as a feature, not a bug: the duality of memory and implications for ubiquitous computing. *CoDesign* 2, 1 (2006), 3–15. arXiv:https://doi.org/10.1080/15710880600608230 doi:10.1080/15710880600608230
- [6] Jakob E. Bardram. 2005. Activity-based computing: support for mobility and collaboration in ubiquitous computing. *Personal and Ubiquitous Computing* 9, 5 (Jan. 2005), 312–322. doi:10.1007/s00779-004-0335-2
- [7] Victoria Bellotti and Keith Edwards. 2001. Intelligibility and Accountability: Human Considerations in Context-Aware Systems. *Human-Computer Interaction* 16, 2-4 (2001), 193–212. arXiv:https://doi.org/10.1207/S15327051HCI16234_05 doi:10.1207/S15327051HCI16234_05
- [8] Ganesh D. Bhatt. 2001. Knowledge management in Organizations: Examining the Interaction Between Technologies, Techniques, and People. *Journal of Knowledge Management* 5, 1 (2001), 68–75.
- [9] Andrea Bunt, Giuseppe Carenini, and Cristina Conati. 2007. Adaptive content presentation for the web. In *The adaptive web: Methods and strategies of web personalization*. Springer, 409–432.
- [10] Rob Comber and Elina Eriksson. 2023. Computing as Ecocide. In *Ninth Computing within Limits 2023* (Online). LIMITS. doi:10.21428/bf6fb269.9fcd0c0c https://limits.pubpub.org/pub/a8h46wqy.
- [11] Maximilian Dauner and Gudrun Socher. 2025. Energy costs of communicating with AI. *Frontiers in Communication* Volume 10 - 2025 (2025). doi:10.3389/fcomm.2025.1572947
- [12] Paul Dourish. 2004. What we talk about when we talk about context. *Personal and ubiquitous computing* 8 (2004), 19–30. doi:10.1007/s00779-003-0253-8
- [13] Ian Drosos, Advait Sarkar, Xiaotong Xu, Carina Negreanu, Sean Rintel, and Lev Tankelevitch. 2024. “It’s like a rubber duck that talks back”: Understanding Generative AI-Assisted Data Analysis Workflows through a Participatory Prompting Study. In *Proceedings of the 3rd Annual Meeting of the Symposium on Human-Computer Interaction for Work* (Newcastle upon Tyne, United Kingdom) (CHIWORK ’24). Association for Computing Machinery, New York, NY, USA, Article 16, 21 pages. doi:10.1145/3663384.3663389
- [14] Shiran Dudy, Steven Bedrick, and Bonnie Webber. 2021. Refocusing on Relevance: Personalization in NLG. *Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing* (2021), 5190–5202. doi:10.18653/v1/2021.emnlp-main.421
- [15] Josh Dzieza. 2023. Inside the AI Factory: the humans that make tech seem human. *The Verge* (2023). https://www.theverge.com/features/23764584/ai-artificial-intelligence-data-notation-labor-scale-surge-remotasks-openai-chatbots
- [16] Saul Greenberg. 2001. Context as a Dynamic Construct. *Human-Computer Interaction* 16, 2-4 (2001), 257–268. arXiv:https://doi.org/10.1207/S15327051HCI16234_09 doi:10.1207/S15327051HCI16234_09
- [17] Manuel Hoffmann, Sam Boysel, Frank Nagle, Sida Peng, and Kevin Xu. 2025. *Generative AI and the Nature of Work*. Technical Report 25-021. Harvard Business School. http://dx.doi.org/10.2139/ssrn.5007084.
- [18] Mohammad Hossein Jarrahi, David Askay, Ali Eshraghi, and Preston Smith. 2023. Artificial intelligence and knowledge management: A partnership between human and AI. *Business Horizons* 66, 1 (2023), 87–99. doi:10.1016/j.bushor.2022.03.002
- [19] Blooma John, Zeena Alsamarra'i, and Niki Panteli. 2023. Enhancing Employee Experience in the Era of Hybrid Work: The Case of Microsoft Viva. *IEEE Software* 40, 2 (2023), 70–79. doi:10.1109/MS.2022.3229956
- [20] Majeed Kazemitabaar, Jack Williams, Ian Drosos, Tovi Grossman, Austin Zachary Henley, Carina Negreanu, and Advait Sarkar. 2024. Improving Steering and Verification in AI-Assisted Data Analysis with Interactive Task Decomposition. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology* (Pittsburgh, PA, USA) (UIST ’24). Association for Computing Machinery, New York, NY, USA, Article 92, 19 pages. doi:10.1145/3654777.3676345
- [21] Tae Soo Kim, Yoonjoo Lee, Jamin Shin, Young-Ho Kim, and Juho Kim. 2024. EvalLM: Interactive Evaluation of Large Language Model Prompts on User-Defined Criteria. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (, Honolulu, HI, USA,) (CHI ’24). Association for Computing Machinery, New York, NY, USA, Article 306, 21 pages. doi:10.1145/3613904.3642216

- [22] Yoonsu Kim, Jueon Lee, Seoyoung Kim, Jaehyuk Park, and Juho Kim. 2024. Understanding Users' Dissatisfaction with ChatGPT Responses: Types, Resolving Tactics, and the Effect of Knowledge Level. In *Proceedings of the 29th International Conference on Intelligent User Interfaces* (Greenville, SC, USA) (IUI '24). Association for Computing Machinery, New York, NY, USA, 385–404. doi:10.1145/3640543.3645148
- [23] Susan Leigh Star Laura J. Neumann. 1996. Making Infrastructure: The Dream of a Common Language. In *Proceedings of the 4th Biennial Participatory Design Conference 1996* (Cambridge, Massachusetts, USA) (PDC '96). Computer Professionals for Social Responsibility, Palo Alto, CA, USA, 231–240. <https://ojs.ruc.dk/index.php/pdc/article/view/153>
- [24] Sasha Luccioni, Yacine Jernite, and Emma Strubell. 2024. Power Hungry Processing: Watts Driving the Cost of AI Deployment?. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency* (Rio de Janeiro, Brazil) (FAccT '24). Association for Computing Machinery, New York, NY, USA, 85–99. doi:10.1145/3630106.3658542
- [25] Olivera Marjanovic, Dubravka Cezec-Kecmanovic, and Richard Vidgen. 2018. Algorithmic Pollution: Understanding and Responding to Negative Consequences of Algorithmic Decision-Making. In *Living with Monsters? Social Implications of Algorithmic Phenomena, Hybrid Agency, and the Performativity of Technology*, Ulrike Schultze, Margunn Aanestad, Magnus Mähring, Carsten Østerlund, and Kai Riemer (Eds.). Springer International Publishing, Cham, 31–47.
- [26] Microsoft. 2025. Microsoft 365 Copilot overview. <https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-overview>. Accessed: 03 Jan 2025.
- [27] Harri Oinas-Kukkonen. 2018. Personalization Myopia: A Viewpoint to True Personalization of Information Systems. In *Proceedings of the 22nd International Academic Mindtrek Conference* (Tampere, Finland) (Mindtrek '18). Association for Computing Machinery, New York, NY, USA, 88–91. doi:10.1145/3275116.3275121
- [28] Wanda J. Orlikowski. 2002. Knowing in Practice: Enacting a Collective Capability in Distributed Organizing. *Organization Science* 13, 3 (2002), 249–273. <http://www.jstor.org/stable/3086020>
- [29] Taemin Kim Park. 1993. The nature of relevance in information retrieval: An empirical study. *The library quarterly* 63, 3 (1993), 318–351.
- [30] Linda Schamber, Michael B. Eisenberg, and Michael S. Nilan. 1990. A re-examination of relevance: toward a dynamic, situational definition. *Information Processing & Management* 26, 6 (1990), 755–776. doi:10.1016/0306-4573(90)90050-C
- [31] Kjeld Schmidt and Liam Bannon. 1992. Taking CSCW seriously: Supporting articulation work. *Computer Supported Cooperative Work (CSCW)* 1 (March 1992), 7–40. doi:10.1007/bf00752449
- [32] Chirag Shah and Emily M. Bender. 2024. Envisioning Information Access Systems: What Makes for Good Tools and a Healthy Web? *ACM Trans. Web* 18, 3, Article 33 (April 2024), 24 pages. doi:10.1145/3649468
- [33] Susan Leigh Star and Karen Ruhleder. 1994. Steps towards an ecology of infrastructure: complex problems in design and access for large-scale collaborative systems. In *Proceedings of the 1994 ACM Conference on Computer Supported Cooperative Work* (Chapel Hill, North Carolina, USA) (CSCW '94). Association for Computing Machinery, New York, NY, USA, 253–264. doi:10.1145/192844.193021
- [34] Miriam Sturdee, Joseph Lindley, Conor Linehan, Chris Elsdon, Neha Kumar, Tawanna Dillahunt, Regan Mandryk, and John Vines. 2021. Consequences, Schmonsequences! Considering the Future as Part of Publication and Peer Review in Computing Research. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI EA '21). Association for Computing Machinery, New York, NY, USA, Article 95, 4 pages. doi:10.1145/3411763.3441330
- [35] Na Sun and Donald Kalar. 2025. Gemini at Work: Knowledge Workers' Perceptions and Assessment of Productivity Gains. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference (DIS '25)*. Association for Computing Machinery, New York, NY, USA, 3681–3695. doi:10.1145/3715336.3735679
- [36] Lev Tankelevitch, Viktor Kewenig, Auste Simkute, Ava Elizabeth Scott, Advait Sarkar, Abigail Sellen, and Sean Rintel. 2024. The Metacognitive Demands and Opportunities of Generative AI. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 680, 24 pages. doi:10.1145/3613904.3642902
- [37] Monideepa Tarafdar, Xinru Page, and Marco Marabelli. 2023. Algorithms as co-workers: Human algorithm role interactions in algorithmic work. *Information Systems Journal* 33, 2 (2023), 232–267. doi:10.1111/isj.12389
- [38] Ben Lee Taylor. 2023. Long hours and low wages: the human labour powering AI's development. *The Conversation* (2023). <https://theconversation.com/long-hours-and-low-wages-the-human-labour-powering-ais-development-217038>
- [39] Jaime Teevan, Susan T. Dumais, and Eric Horvitz. 2010. Potential for personalization. *ACM Trans. Comput.-Hum. Interact.* 17, 1, Article 4 (April 2010), 31 pages. doi:10.1145/1721831.1721835
- [40] Elaine G. Toms. 2011. Task-based Information Searching and Retrieval. In *Interactive Information Seeking, Behaviour and Retrieval*, Ian Ruthven and Diane Kelly (Eds.). Facet Publishing, London, United Kingdom, Chapter 3, 43–59. ISBN: 978-1-85604-707-4.
- [41] Tung Vuong, Salvatore Andolina, Giulio Jacucci, and Tuukka Ruotsalo. 2021. Does More Context Help? Effects of Context Window and Application Source on Retrieval Performance. *ACM Trans. Inf. Syst.* 40, 2, Article 39 (Sept. 2021), 40 pages. doi:10.1145/3474055
- [42] Tung Vuong and Tuukka Ruotsalo. 2024. Predicting Representations of Information Needs from Digital Activity Context. *ACM Trans. Inf. Syst.* 42, 4, Article 95 (Feb. 2024), 29 pages. doi:10.1145/3639819
- [43] Ryan W. White and Chirag Shah (Eds.). 2024. *Information Access in the Era of Generative AI*. Springer Cham. doi:10.1007/978-3-031-73147-1
- [44] Denise Wilkins, Siân Lindley, Max Meijer, Richard Banks, and Britta Burlin. 2020. Designing AI Systems That Make Organizational Knowledge Actionable. *Interactions* 27, 6 (nov 2020), 72–75. doi:10.1145/3424690
- [45] John Winn, Matteo Venanzi, Tom Minka, Ivan Korostelev, John Guiver, Elena Pochernina, Pavel Myshkov, Alex Spengler, Denise Wilkins, Siân Lindley, Richard Banks, Sam Webster, and Yordan Zaykov. 2021. Enterprise Alexandria: Online High-Precision Enterprise Knowledge Base Construction with Typed Entities. In *Automated Knowledge Base Construction*. <https://www.microsoft.com/en-us/research/publication/enterprise-alexandria-online-high-precision-enterprise-knowledge-base-construction-with-typed-entities/>
- [46] Allison Woodruff, Renee Shelby, Patrick Gage Kelley, Steven Rousso-Schindler, Jamila Smith-Loud, and Lauren Wilcox. 2024. How Knowledge Workers Think Generative AI Will (Not) Transform Their Industries. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 641, 26 pages. doi:10.1145/3613904.3642700
- [47] Ryan Yen, Nicole Sultanum, and Jian Zhao. 2024. To Search or To Gen? Exploring the Synergy between Generative AI and Web Search in Programming. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA '24). Association for Computing Machinery, New York, NY, USA, Article 327, 8 pages. doi:10.1145/3613905.3650867
- [48] Danni Yu. 2023. The 'AI divide' between the Global North and the Global South. *World Economic Forum* (2023). <https://www.weforum.org/stories/2023/01/davos23-ai-divide-global-north-global-south/>
- [49] Jingtao Zhan, Qingyao Ai, Yiqun Liu, Jia Chen, and Shaoping Ma. 2024. Capability-aware Prompt Reformulation Learning for Text-to-Image Generation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Washington DC, USA) (SIGIR '24). Association for Computing Machinery, New York, NY, USA, 2145–2155. doi:10.1145/3626772.3657787