

"If You're Very Clever, No One Knows You've Used It": The Social Dynamics of Developing Generative AI Literacy in the Workplace

Qing (Nancy) Xia
University College London
London, United Kingdom
nancy.xia.18@ucl.ac.uk

Marios Constantinides*
CYENS Centre of Excellence
Nicosia, Cyprus
marios.constantinides@cyens.org.cy

Advait Sarkar†
Microsoft Research
Cambridge, United Kingdom
advait@microsoft.com

Duncan P. Brumby
University College London
London, United Kingdom
d.brumby@ucl.ac.uk

Anna L. Cox‡
University College London
London, United Kingdom
anna.cox@ucl.ac.uk

Abstract

Generative AI (GenAI) tools are rapidly transforming knowledge work, making AI literacy a critical priority for organizations. However, research on AI literacy lacks empirical insight into how knowledge workers' beliefs around GenAI literacy are shaped by the social dynamics of the workplace, and how workers learn to apply GenAI tools in these environments. To address this gap, we conducted in-depth interviews with 19 knowledge workers across multiple sectors to examine how they develop GenAI competencies in real-world professional contexts. We found that, while knowledge sharing from colleagues supported learning, the ability to remove cues indicating GenAI use was perceived as validation of domain expertise. These behaviours ultimately reduced opportunities for learning via knowledge sharing and undermined transparency. To advance workplace AI literacy, we argue for fostering open dialogue, increasing visibility of user-generated knowledge, and greater emphasis on the benefits of collaborative learning for navigating rapid technological developments.

CCS Concepts

• Human-centered computing → HCI theory, concepts and models; Empirical studies in HCI.

Keywords

Generative AI, AI literacy, Knowledge work, Competences

ACM Reference Format:

Qing (Nancy) Xia, Marios Constantinides, Advait Sarkar, Duncan P. Brumby, and Anna L. Cox. 2026. "If You're Very Clever, No One Knows You've Used It": The Social Dynamics of Developing Generative AI Literacy in the Workplace. In *Proceedings of the 5th Annual Symposium on Human-Computer Interaction for Work (CHIWORK '26)*, June 22–25, 2026, Linz, Austria. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3808045.3808050>

*Also affiliated with University of Cyprus, Cyprus and University College London, UK.

†Also affiliated with University of Cambridge, University College London, UK

‡Also affiliated with University of Bristol, UK and CYENS Centre of Excellence, Cyprus



This work is licensed under a Creative Commons Attribution 4.0 International License. *CHIWORK '26, Linz, Austria*

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

ACM ISBN 979-8-4007-2429-9/26/06

<https://doi.org/10.1145/3808045.3808050>

1 Introduction

Generative Artificial Intelligence (GenAI) tools are rapidly redefining modern knowledge work [23, 36, 37, 45]. Adoption rates have increased annually since 2023 [105, 106]. Around 78% of 31,000 knowledge workers surveyed in Microsoft and LinkedIn's 2024 Annual Report of Work Trends Index reported bringing their own AI tools to work [72]. As a result, there has been increased interest in reskilling existing employees to newly automated workflows, and ensuring responsible usage by increasing users' awareness of the social and ethical implications of GenAI use [69, 75, 105, 106, 139].

Establishing AI literacy is now essential. AI literacy describes the competencies necessary for individuals to understand, interact with, and critically evaluate AI technologies [6, 66, 86]. Important competencies often include understanding the technical underpinnings of AI [18, 66]; interaction techniques (e.g., prompt design [55]); variations of learning, critical thinking, reflection, and metacognitive skills [28, 58, 68, 74, 86]; and socio-technical competencies, such as understanding the role of human actors in AI use through the lens of social and ethical impacts (e.g., [8, 18, 86]). Lower levels of AI literacy are consistently linked with a range of problems, including a tendency to over-rely on AI tools [88, 117, 136], be more susceptible to hallucinations and biases [83], lower job performance compared to peers with higher AI literacy [65], and lower employability [89].

Significant effort has been directed towards defining and designing for student AI literacy in formal educational environments [22, 74, 131]. However, professional users often rely on more unstructured methods of learning (e.g., informal knowledge sharing [51, 59, 99]) to develop highly personalised understandings of how and when different tools should be applied in their work, that can differ from those emphasised in existing AI literacy frameworks. Indeed, current GenAI adoption among employees often occurs independently of organisational guidance and strategy: surveys indicate that 53% of employees do not disclose their GenAI use to their employers [72], while McKinsey & Company's 2025 analysis showed that C-suite leaders significantly underestimated the extent of such adoption in day-to-day work [69]. Many professionals are therefore constructing their own understanding of GenAI competencies and learning strategies, shaped less by organisational guidance or formal literacy frameworks, and more by personal priorities and their wider social context [46].

As the attitudes and beliefs of today's workforce will shape future norms of interaction and learning, it is vital to understand how knowledge workers develop GenAI competencies in practice. Specifically, we explore perspectives that elaborate and contextualise workers' beliefs and learning strategies for GenAI competencies, focusing in particular on extending understanding of overlooked socio-technical competencies (see Table 1). We aim to address two main Research Questions (RQs):

RQ₁: What competencies do knowledge workers value when using GenAI tools in the workplace?

RQ₂: Which strategies do knowledge workers consider essential to developing these competencies?

In answering these questions, we make three main contributions:

- (1) We conducted semi-structured interviews with 19 knowledge workers from diverse professions, allowing us to develop an empirically grounded account of how GenAI literacy is practised and negotiated within our sample.
- (2) We showed that, while workers benefitted from colleagues sharing knowledge, the act of critiquing and obscuring their own GenAI usage was also valuable to them. Importantly, while prior work has interpreted acts of GenAI hiding as rooted in shame [137], or a desire to elevate one's status [98], we found that users can also view successful erasure of telltale signs of GenAI outputs as a positive indication of one's professional expertise. However, these behaviours have negative impacts on further opportunities to learn via knowledge sharing, and promote cultures that lack transparency (Section 4).
- (3) We extend existing AI and GenAI literacy frameworks by highlighting the need to consider human stakeholders both as a resource for learning and as a source of social influence on usage. Building on this perspective, we highlight implications for HCI design to support open dialogue, disclosure, and sustainable knowledge sharing in the workplace (Section 5).

2 Related Work

Our work builds on several lines of research that may be grouped into three areas: (1) GenAI in knowledge work (Section 2.1); (2) developing digital skills in the workplace (Section 2.2); and (3) defining AI and GenAI literacy (Section 2.3).

2.1 GenAI in Knowledge Work

Knowledge work describes the complex and non-routine labour that requires tailored problem-solving skills and high levels of expertise [31, 34, 113]. Knowledge workers in various sectors (e.g., researchers, consultants, creatives, IT professionals [32, 113]) comprise a significant part of the world economy. In recent years, there has been an increase in interest and uptake in the use of GenAI tools to knowledge work [23, 36, 37, 45]. We define GenAI as end-user tools that draw on deep learning approaches and large amounts of training data to create outputs in response to user prompts [11, 23, 97]. The integration of GenAI tools into knowledge work is expected to bring benefits to productivity, but also challenges [103]. GenAI tools can conduct many core components of knowledge

work tasks [13, 93, 116], such as searching and collating information [107, 132, 138], communication [3, 112], or generating tailored solutions in problem-solving [12, 94, 120]. However, recent work suggests that the productivity gains from automating these tasks may be offset by demands for new skill-sets which arise from using GenAI [26]. Multiple studies have highlighted that the use of GenAI tools presents a shift in users' labour from production to supervision [57, 104, 126]. This shift introduces new tasks to pre-existing workflows, such as refining prompts or adapting outputs (i.e., task stewardship) [30, 55, 57, 135], but also require users to make informed decisions as to when and what types of tasks could be delegated to GenAI [48, 54], termed 'critical integration' [95]. However, this introduces new challenges. Tankelevitch et al. [115] highlight how adoption of GenAI introduces significant metacognitive effort by requiring users to clearly articulate and refine their goals when prompting [55, 135]. Reliance on low quality generated GenAI outputs can also threaten the need for human expertise and risk deskilling industries, negatively impact critical thinking and cognition [57], or promote practice that lose the human characteristics (e.g., empathy) for interpersonal work [54, 126]. As GenAI reshapes the skills and cognitive demands placed on knowledge workers, it becomes increasingly important to re-evaluate which competencies and learning strategies users prioritise in these evolving work environments, and their motivations for doing so.

2.2 Developing Technical Skills in the Workplace

While traditional educational models in school settings are characterised by structured, theoretical progression and clearly defined performance indicators, learning in the workplace is often more participatory, and dependent on communities of practice [123], where each worker is expected to play an active role in creating and sharing knowledge [40, 119]. The social context and priorities of the workplace play an important role in how learning takes place.

Previous work exploring the development of technical skills (e.g., software use, end-user programming) in the workplace highlights that learning is often shaped by external factors such as time pressures, competing priorities, and a pressure to maintain productivity [59, 102]. As a result, workers tend to rely on strategies that are perceived as readily available and directly relevant, such as informal help-seeking, knowledge sharing with colleagues, or turning to search engines [51, 59, 118]. Furthermore, an individual's technological competence is often closely intertwined with their professional identity and value [77, 96]. As a result, users' personal interests and values regarding what technical skills needed are also influenced by the social norms and values of the environment [46, 77–80]. For example, recent studies of spreadsheet users highlight that users may alter the functionalities (e.g., formula-writing vs. using Macros) or visual layout of a shared spreadsheet to conform to the needs of the recipient and to avoid incurring negative or unprofessional social cues (e.g., messy presentations) that may be evoked by the way the spreadsheet is presented [21, 110, 127].

Concerns about social scrutiny and judgement based on one's technical competence has been a pervasive barrier to workplace knowledge sharing in a software learning context [42, 87], but have become even more relevant in the context of GenAI due to trends of

stigmatisation against GenAI users. Recent works have highlighted that disclosure of GenAI use is often accompanied with undesirable social consequences, such as lower perceived credibility, quality of work, and lowered trust [90, 101]. This social environment has led some researchers to explore the issue of AI hiding and shame, a behaviour whereby individuals downplay or intentionally remove evidence of AI use to avoid the associated social penalty incurred through admitting its usage [61, 98, 137]. For example, once people became aware that GenAI outputs were often biased with specific word choices (e.g., 'delve'), there were notable decreases in the word's occurrence as users sought to disassociate their work from GenAI use [41]. However, while AI hiding is seen as a behaviour that counters the overall goal of establishing transparency in responsible GenAI usage [101], its impact on the knowledge sharing dynamics that are important for technical skill development in the workplace is not well understood. As knowledge sharing is a key method for retaining and maintaining organisational knowledge and technical skillsets among employees long-term [25, 59, 60, 108], understanding the impacts of AI hiding through the lens of knowledge sharing can have important implications for considering how AI literacy could be sustainably supported over time.

2.3 Defining AI and GenAI Literacy

The traditional concept of 'literacy' is about ensuring that members of the population are equipped with the foundational skills needed to understand, express, and communicate using language. In recent years, there has been growing interest across the fields of computer education and HCI in exploring the concept of literacy in relation to AI technologies [6, 28], including but extending beyond GenAI to include robotics, machine learning etc. [66, 86]. This research is largely divided into two strands. The first, rooted in the field of education, primarily focuses on defining or measuring AI literacy [19, 46, 66, 86, 122]. The second, which has been of greater focus in the HCI community, is interested predominantly in supporting the development of AI literacy across a diverse range of users [17, 49, 53, 56, 130]. More recently, works have also drawn distinctions between AI and GenAI literacy. This is due to the widespread popularity and accessibility of GenAI tools among non-specialist users, the need for more specific skills such as prompt optimisation when interacting with GenAI tools, as well as the cross-domain applicability of GenAI tools requiring more careful consideration of how and when such technology should be applied in different contexts [8, 64].

We observe three main dimensions in which proposed competencies may be clustered: 1) technological competencies, describing the technical knowledge needed to understand and interact with AI; 2) learning and reflective competencies, which covers more general metacognitive and thinking skills of reflection, evaluation, and communication; and 3) socio-technical competencies, which involves understanding the role of human stakeholders throughout the process of AI use. Table 1 summarises AI and GenAI literacy frameworks across the three dimensions, and highlights how each defines competencies for adult and general audiences. Due to the limited number of GenAI literacy frameworks and the significant overlaps between GenAI and AI literacy frameworks, we review both categories of framework together, and annotate competencies that are specific to GenAI.

While socio-technical competencies play an important role for AI literacy, there is a significant focus on educating users on normative ethical implications of GenAI use. However, as stated in Section 2.2, users' personal interests and values regarding different technical skills are closely shaped by their social norms, organisational contexts, and a desire to support the presentation of their own professional identity [77, 79]. Additionally, only a few frameworks identify the need to navigate social and collaborative contexts in AI use [46, 68, 86], and they often overlook how such dynamics shape sustained learning.

In summary, our review highlights that our understanding of the competencies needed to use GenAI successfully in knowledge work is still evolving. We also find that current AI literacy frameworks largely overlook the social dynamics of AI shaming, disclosure, and hiding, and in particular its implications for unstructured learning within the workplace. We aim to fill this gap by providing a deeper exploration of the beliefs and experiences of knowledge workers who are continuously learning to use GenAI tools while managing the social dynamics of the workplace.

3 Method

3.1 Participants

During participant recruitment, we first identified sectors, which, based on previous economic analyses, were adopting generative AI tools at rapid rates [23, 37, 69, 106]. We then identified three core characteristics of 'knowledge work' and recruited participants belonging to sectors that met those characteristics. We recruited from sectors where expertise is developed primarily through formal education or training [34] (i.e., IT, Data Science, Law), sectors where tasks involved non-routine or creative problem-solving that require contextual sensitivity to meet stakeholders' needs [31] (i.e., Management and Human Resources, Creative industries), and sectors which specialised in knowledge creation and dissemination [32] (i.e., Research and Development). Incorporating diverse perspectives from different sectors is important for ensuring rigour in qualitative research, as our goal is not to produce a generalisable account of human behaviour, but to ensure that we extend existing interpretations of the research subject [15].

In order to better access participants from these different occupational roles, we used the online recruitment platform Prolific. Prior evaluations has demonstrated that data quality from participants on Prolific, evaluated through user attention, comprehension, honesty, and care with answering questions, was consistently higher than comparable platforms [4, 33, 82]. One of the advantages of using Prolific was that it allowed finer control over participants' demographics and attributes, allowing recruitment to be targeted. We used both Prolific's own in-built filters, based on participants' self-reported demographics when registering for a Prolific account, as well as our own screening survey to ensure participants met our recruitment criteria, namely, participants should: have occupations belonging to one of the sectors identified above; use GenAI tools at least once a week and primarily for work-related (as opposed to leisure or personal) purposes; be employed in a full- or part-time capacity in their specified sector and role; have more than 1 year of experience in their current role; and are not working in organisations which explicitly prohibited the use of GenAI tools in their

Table 1: List of technological, socio-technical, and learning and reflective dimensions with corresponding interpretations of competencies across AI and GenAI literacy frameworks.

Dimension	Interpretation	References
Technological competencies	Understanding what AI/GenAI is, what it is capable of, how it works, and its limitations	[8, 18, 19, 28, 46, 58, 65, 66, 74, 85, 86]
	Ability to apply AI tools to solve problems, complete tasks, and enhance productivity and work outputs	[8, 18, 19, 28, 46, 58, 65, 66, 74, 85, 86, 133]
	Creating and designing AI models, algorithms, and tools	[8, 18, 85]
	(<i>GenAI-specific</i>) Skill in prompt engineering, articulation and proficiency in communication with GenAI tools	[8, 58, 65]
Learning and reflective competencies	Understanding how AI may be perceived, and its impact on users' experiences, social interactions with AI, how accessible and learnable AI may appear to be	[66]
	Critical thinking and evaluation of the outputs of AI or GenAI models	[58, 65, 74, 86]
	Metacognitive ability to reflect on, identify, address and manage one's own needs (e.g., what information is needed, how to access it); and how external factors (e.g., AI) may influence the thinking process	[18, 28, 46, 68, 86, 133]
	Ability to think creatively about using AI, to translate innovative approaches into practical outcomes and ideas	[46, 58, 65, 68]
	Ability to engage in continuous learning to adapt to developments in the technology	[8, 18, 19, 68, 86]
	Foundational literacy skills in reading, writing, arithmetic; basic understanding of how to understand and communicate through symbols, text, and numbers	[133]
	Regulation of emotions when interacting with AI	[18]
	(<i>GenAI-specific</i>) Ability to detect GenAI-generated content	[8]
Socio-technical competencies	Consideration of the ethical and legal implications of AI use e.g., equality, privacy, transparency, accountability, impact on employment, environmental impacts, trust, safety	[8, 18, 28, 46, 58, 65, 66, 74, 85, 86]
	Understanding the role of human actors in human-AI collaboration, including human-specific advantages and disadvantages	[19, 68, 85, 86]
	Understanding the AI impact on social norms, relations, and collaboration with other human stakeholders	[46, 68, 86]

work, as this may constrain the types of learning resources they can access, which was an important aspect of the research question. However, we allowed participants to participate if their organisation had no clear GenAI policies, or if participants themselves were uncertain of the policies, as this allowed us to explore how our participants navigated and developed their own understanding of GenAI competencies independent of top-down organisational guidance. These criteria ensured that our participants had sufficient familiarity with their specific work environment, domain, and had

the opportunity to access the relevant GenAI tools and learning resources to be able to discuss their experiences, perceptions, and learning strategies in the interview.

We initially identified 25 eligible participants for interviews. Six participants (not included in the table) were excluded. Four dropped out of the study before the interview or experienced technical issues during the interview, and two were excluded post-interview because they used tools which we could not reliably verify as using GenAI models. Our final data-set consisted of semi-structured

interviews with 19 knowledge workers collected from December 4th 2024 to January 17th 2025, which meets the standard sample size expected for the field [16].

Table 2 outlines the demographics of the participants, their job title, organisation size, and experience. Table 3 outlines the types of GenAI tools used by each participant, and the types of tasks which the GenAI tools are used for.

3.2 Interview Protocol and Data Collection

We developed a semi-structured interview protocol to facilitate data collection. As our participant sample spanned a wide range of sectors, we ensured mutual understanding by establishing at the beginning of interviews: (1) the occupation and professional responsibilities of the participant; 2) their understanding of the term ‘generative AI’; and 3) how participants currently use GenAI tools in their work. This provided opportunities for the interviewer to ask for clarifications regarding participants’ backgrounds, informally evaluate participants’ technical expertise and to adjust their terminology appropriately to ensure communication, and to guide conversations to focus on GenAI tools if they struggled to distinguish between similar technological products (e.g., generative vs. rule-based Chatbots).

During interviews, we prioritised flexibility, open questions, and following up on participants’ statements in order to better capture participants’ individual perspectives, experiences, and beliefs. Prior to data collection, the protocol was also piloted with two participants (one PhD researcher, one creative copywriter) to evaluate the terminology used and to ensure that the questions asked prompted relevant interpretations and discussion. The main topics of inquiry explored in the study involved:

- What competencies do participants perceive to be most important to using GenAI in their work?
- How did participants develop their current practices and workflows for using GenAI tools?
- What resources or strategies do participants use to help them learn to use GenAI tools, and why?
- How do users communicate about their skills and knowledge of GenAI tools within their workplace?

Interviews were conducted online by the first author, using Microsoft Teams. Eligible participants were identified via the initial recruitment survey distributed on Prolific, and invited to book a suitable timeslot for the interview. The interviews recorded and transcribed using Microsoft Teams’ built-in functionality. Transcriptions were then manually corrected in accordance to video recordings. Interviews typically lasted between 50-60 minutes. Ethics for conducting the study was obtained and approved by the University College London internal ethics committee. All participants were informed of the purpose of the study and consented to participate. To accommodate for participants’ needs [5], participants were invited to book timeslots for the interview according to their own availability. All participants were compensated £0.70 for completing the 2-minute screening survey, regardless of whether or not they also completed the interview. All interview participants, including those whose data were ultimately excluded post-data collection, were compensated at £13 an hour, proportional to the amount of time taken in their respective interviews.

3.3 Analysis

3.3.1 Positionality. Our research team consisted of five members with expertise in research with Human-Computer Interaction, responsible GenAI, software learnability, and technology use in the workplace. Our overall approach to collecting, analysing, and presenting the data is rooted in critical realism [14, 84], meaning we recognise that both our interpretation as well as participants’ own articulations reflect their subjective experiences and beliefs about GenAI competencies and the resources necessary to support their learning [44], rather than reflecting an objective or concrete reality. We adopt this position as it allows us to focus explicitly on understanding participants’ personal experiences and values. This is important for addressing the primary goals of this paper, which is to develop alternative ways of thinking, clarifying, and extending existing frameworks of AI literacy and competencies.

Due to the nature of critical realism, it is also important for transparency and rigour to disclose our authorial positionality. The first author, who collected the data and led the analysis process, is an academic researcher with a primary interest in organisational learning. They are guided by the view that learning is an inherently social process, shaped and developed through interactions with a wider community [123]. As such, our research process focused on relevant competencies in GenAI, considering not only the technical knowledge of GenAI interactions or understanding of GenAI models, but also on the skills surrounding GenAI use that are important for users’ social outcomes, such as reputation and maintaining positive relationships [1, 46, 86, 122]. To support reflexivity and critical reflection, the analysis process was supported by two co-authors with expertise in GenAI research, who offered alternative perspectives for the first author. Two additional co-authors with expertise in workplace research further provided reflections over the main themes during write-up. In addition, the first author, who was the only one to interact directly with participants, had no prior experience with any non-academic sectors participants were affiliated with. This allowed for consistency during data collection when questioning and establishing an understanding of participants’ work contexts and their GenAI usage.

3.3.2 Thematic analysis. We conducted thematic analysis using a template approach [52]. An initial framework, i.e., a ‘template’, is applied to scaffold the initial coding and interpretation of the data, but new categories may be added or merged during coding to reflect researchers’ evolving understanding of the data. This method is particularly suitable for identifying nuances where interpretations of the data diverge from pre-existing frameworks. While we acknowledge that there is no widely agreed definition of the competencies necessary to become AI literate, and that the importance of different competencies vary depending on researchers’ domain and pedagogy [6, 68], the aim of our analysis was to provide insights into the underlying reasoning and values behind how knowledge workers conceptualise and develop GenAI competencies.

We applied different theoretical frameworks to organise the initial analysis of our data, corresponding to each research questions. For RQ1, we used two AI literacy frameworks to provide an initial mapping of how participants conceptualised GenAI competencies. The first framework was developed by Pinski and Benlian [86], and was chosen because: (1) it specifies competencies shaped by AI’s

Table 2: Demographics of the 19 interview participants, including the participant’s ID, age range and gender, job context, organisation size, and domain experience. Participants represented diverse sectors such as academia, law, IT, creative industries, finance, and human resources, with experience levels ranging from 1–4 years to over 10 years. Participant IDs are not in order as six initially eligible participants were excluded from the final data-set.

ID	Age (Gender)	Job Context	Organisation Size	Domain Experience
P1	45-54(M)	Astrophysicist and academic researcher	Large (250+ employees)	1-4 years
P2	25-34(F)	Bid writer in marketing agency	Small (10-50 employees)	1-4 years
P4	25-34(M)	Social media marketing and communications manager	Large (50-250 employees)	1-4 years
P6	45-54(F)	Data analyst in mental health charity	Large (250+ employees)	5-9 years
P7	18-24(M)	Data analyst in insurance company	Large (250+ employees)	1-4 years
P9	35-44(M)	Creative copywriter for consumer technology company	Large (250+ employees)	5-9 years
P10	25-34(F)	Cybersecurity engineer	Medium (50-250 employees)	5-9 years
P12	55-64(M)	Software and solutions architect specialising in medical systems	Large (250+ employees)	10+ years
P13	55-64(M)	University lecturer and researcher	Large (250+ employees)	10+ years
P14	25-34(M)	Social media manager and assistant HR	Small (10-50 employees)	5-9 years
P15	45-54(F)	Research professor in digital business	Large (250+ employees)	10+ years
P17	25-34(F)	Trainee solicitor	Medium (50-250 employees)	1-4 years
P18	45-54(F)	HR officer in agricultural institution	Small (10-50 employees)	10+ years
P19	25-34(M)	Sales coordinator, analyst, and distributor for a transport logistics company	Small (10-50 employees)	1-4 years
P20	25-34(F)	Senior accountant	Medium (50-250 employees)	1-4 years
P21	25-34(M)	Software engineer	Medium (50-250 employees)	1-4 years
P22	25-34(M)	Business relationships manager	Large (250+ employees)	5-9 years
P23	25-34(M)	Senior architect	Medium (50-250 employees)	1-4 years
P25	25-34(F)	Marketing coordinator for a financial software company	Medium (50-250 employees)	5-9 years

unique features (e.g., inscrutability, autonomy, and its ability to learn) which challenge the assumptions of conventional computer learning; (2) it provides a socio-technical perspective that goes beyond technical knowledge to include the social skills needed to navigate AI use with other stakeholders; and (3) it was developed based on a recent literature review of AI literacy. The second framework was the GenAI literacy framework proposed by Annapureddy et al. [8], which, to our knowledge, is one of the only attempts to date to provide a systematic overview of the competencies unique to GenAI. Related categories from both frameworks were merged together for the initial template.

To address RQ2 and understand the general facilitators and barriers which affect users’ choice of learning resources and strategies, we used the COM-B model of behaviour to form the initial template for analysis. This model suggests that observable behaviours (in our case, strategies for learning such as trial-and-error) manifest when three core components of behaviour align – capability, opportunity, and motivation. Each component can also influence the other [71]. Though typically applied to help researchers map out and develop their understanding of a single behaviour with the goal of providing prescriptive recommendations, we applied this model to provide a comprehensive overview of the external and internal factors driving users’ choices of GenAI learning strategies.

The first author conducted open coding over the interview datasets, which were then organised according to the three theoretical frameworks outlined above. Particular focus was placed on categories where data diverged from, clarified, or identified areas of overlap, between the two AI literacy frameworks used. We established rigour in the analysis by ensuring that the first author engaged in principles of reflexivity, and by incorporating the perspectives of the second and third authors throughout analysis [14]. Both the second and third author regularly reviewed quotes, code, and theme descriptions as they were being developed. This involved interrogating the overall clarity of themes, the first authors’ interpretations of the relationships between contributing factors, and the evidence selected to support claims. This ensured critical reflection throughout the coding and writing process via multiple alternative perspectives.

4 Results

In answer to RQ1, we identified two core competencies which our participants described as important for using GenAI in practice: (1) a technological competency, namely, the awareness of potential GenAI use cases (see Section 4.1); and (2) a socio-technical competency, namely, the ability to remove cues of GenAI usage via concealment and critique (detailed in Section 4.3) as a means to manage quality of work, as well as one’s own presentation and

Table 3: Participants’ reported use of GenAI tools and their use cases in the workplace. Tools included ChatGPT, Copilot, Claude, Gemini, DALL-E, and others. These tools were used for a range of tasks such as content generation, coding, debugging, writing support, design, communication, and knowledge management.

ID	GenAI Tool(s)	GenAI Use Case
P1	ChatGPT	Generating codes for writing software
P2	Notion AI, ChatGPT	Generating initial drafts and refining writing
P4	ChatGPT, Quillbot, Gemini	Generating content and ideas for social media, refining writing
P6	ChatGPT	Generating Visual Basic Advance code and Excel formulas, providing support and solutions to problems in analysis
P7	ChatGPT, PowerBI/Excel Copilot plugin	Providing advice for debugging, structuring data, performing data analysis according to parameters and analysis goals
P9	Copilot, ChatGPT	Generating ideas for content, image generation to support communication with designers
P10	Claude AI, ChatGPT	Generating initial drafts of documentation and policy reports (e.g., exercises to support training on information security)
P12	Copilot	Answering questions, generating code, providing support for debugging, error checking and conducting unit tests on codes
P13	Gemini, Copilot, ChatGPT	Generating lecture slides, generating summaries of papers for writing papers, writing emails, refining and editing writing
P14	ChatGPT, DALL-E, Amper Music	Generating content for social media, generating code for building websites
P15	ChatGPT	Generating titles for papers, drafting emails, reviewing and generating content ideas for lecture courses
P17	Lexis+ AI	Answering queries and providing support based on precedent case files
P18	Gemini, ChatGPT	Drafting communications and emails, supporting user learning by summarising and teaching agricultural terms when writing investor reports
P19	Claude AI, Lovable.dev, ChatGPT	Creating solutions to support journey mapping and logistics for arranging transport deliveries, generating code to create web platforms for clients, debugging code
P20	ChatGPT, Copilot	Answering queries and providing reminders of accountancy principles to support auditing process
P21	ChatGPT, Gemini, Perplexity, Copilot	Support interpretation of legacy SQL code, conduct online search to return documentation of code libraries, writing emails
P22	ChatGPT, myKnowledge	Answering work-related queries to help with responding to customer queries, proposing solutions
P23	ChatGPT, DALL-E, Claude AI, Gemini	Generating interior designs in different styles, generating Powerpoints for work presentations
P25	Gemini, ChatGPT	Proofreading writing and content for blogs, summarising concepts to support learning about search engine optimisation and marketing-related skills

image to others. In answer to RQ2, we found that these competencies are developed through two corresponding learning strategies: (1) learning from others was vital for developing awareness of potential GenAI use cases (see Section 4.2); and (2) learning through hands-on experience with GenAI tools played an important role in developing users’ ability to detect and mask evidence of GenAI usage, as they become more accustomed to biases and limitations in the GenAI outputs (see Section 4.4).

Taken together, these findings reveal a central tension. While the development of technological competencies, in particular developing awareness of potential AI use cases, benefit from social processes, knowledge exchange, and learning from others, individuals exercising socio-technical competences and considering the social impacts of GenAI use often engage in prevalent practices of concealing and critiquing GenAI usage in ways which constrain opportunities for open knowledge sharing. Importantly, concealment was not always framed negatively as shame. For some, successfully erasing traces of GenAI use was a positive signal of professional expertise. Figure 1 summarises the main themes for each of research question, how they relate to each other, and the contributing factors defining each theme.

4.1 Technological Competency: Awareness of Potential GenAI Use Cases

Discoveries of novel GenAI tools, use cases, and methods of interactions were often described as key moments which redefined participants’ understanding of GenAI technology. P4, describing how her awareness developed over time, stated: *"Before, I just used [ChatGPT] for, let's say, text, right? [...] And I found out, OK, you can use it to generate images, you can use it actually to do analysis."* Awareness helped people feel like they could better understand the impact of GenAI tools, what it meant for the areas of work or the types of tasks that they were interested in, and helped them to access new possibilities for improving their current workflow and practices involving GenAI.

4.1.1 Awareness provides a competitive advantage. Maintaining awareness of GenAI’s potential use cases was important to participants because GenAI tools possessed potential utility across a variety of domains. For participants such as P14, it was empowering to feel that they may be able to harness such technology for their own purposes: *"I know virtually everything now, can be done [...]. You can solve any, most of every problem now, with AI [...]. The next thing was me learning myself. That's what puts me on top of my game."* For some participants, the scope of GenAI’s potential

RQ1: What competencies do knowledge workers value when using GenAI tools in the workplace?

RQ2: Which strategies do knowledge workers consider essential to developing these competencies?




Figure 1: Relationships between competencies and learning strategies identified in the study. Knowledge workers valued awareness of potential GenAI applications and the ability to remove cues indicating GenAI use (detailed in Section 4.3) These competencies were reinforced by two key learning strategies (i.e., learning from others and hands-on experience), while removing cues negatively impacted opportunities for shared learning.

functionalities included not only the technology's current capabilities, but also its anticipated future developments. P19, who closely followed GenAI industry leaders on the social media platform X, explained: "If [industry leaders] say that you[re] build[ing] Sora [video generating AI tool], I'm trying to understand what would be the end goal of Sora? Is it something that could actually replace the movies that we watch?" Participants were motivated to keep up with GenAI developments in order to harness it more quickly and address their current problems, or simply to reduce the perceived threat to their own job and career security by attempting to predict future industry trends. P10 shared: "If care is not taken, children who are 15 years old now [...] they might know how to use this tool, AI, more effectively and efficiently than I do, and before I realise, that could also be a threat to my own job or my career."

However, not all participants were motivated to learn due to career concerns. This was particularly the case when GenAI tools were applied and packaged to users as tools serving very specific or narrow functionalities. One example is P22, whose usage of GenAI primarily consisted of an intelligent internal search engine and knowledge management tool recently implemented in their organisation. He explained: "For learning about something that's already implemented and you're not going to utilise it further than literally just [submitting queries], you're going to learn a lot of areas

which [...] is going to be absolutely pointless [...]. It would be a massive waste of my time."

4.2 Strategy for Learning: Learning from Others

4.2.1 Addressing knowledge gaps. As awareness of potential GenAI use cases was important for increasing users' competitive advantage as knowledge workers, participants were wary of knowledge gaps that may arise due to limitations in their own imagination, interests, and access to information. P15 described how recommended posts on LinkedIn, a professional social networking platform, helped her identify new directions for learning: "On LinkedIn, when I don't look for things, these things are posted, and they'd come to me and I'd become aware of something that I wouldn't be aware of, right?" P12, a software engineer, explained that while his preferred approach towards adopting new technology is to explore and apply trial-and-error until he obtains an ideal outcome, "when you go on your own, you learn the bits that you find interesting. There may be gaps in what you do [...] There could be things there that you're missing."

Incorporating the perspectives of others – regardless of whether their experience aligned with users' domains – was seen as a key way of addressing knowledge gaps. Informal, unstructured knowledge sharing in the form of discussions or casual conversations was a common avenue for learning. Several individuals shared their discoveries of new GenAI tools with colleagues, as P23 explained:

"Sometimes I do teach [my colleagues] examples of AI tools I use [...] from their own perspective, they can also tell me that this [tool] is also better." For some participants, knowledge sharing during the initial adoption of GenAI tools was naturally incorporated into structured weekly meetings: "It was just a natural topic of conversation [...] 'What are people using it for? How useful do they find it?' [...] One of the things that came out of [these meetings] was [using Copilot to conduct] the unit test. It hadn't occurred to me to do that, so I went and looked at that further." (P12).

4.2.2 Accelerate learning and adoption process. For some individuals collective learning and experimentation helped to accelerate a team's mental model of different GenAI tools: "If a tool isn't particularly good at something, then you'd just be going around in circles [...] It's just a lot quicker if there's a few of you trying these things out and then sharing that it worked or not." (P9) While often, participants cited altruistic motivations for sharing GenAI knowledge, in some cases, users strategically share knowledge with target individuals to normalise GenAI usage, as P18 shared: "I find myself most of the time encouraging my supervisor to encourage his bosses to incorporate the use of AI [...] Him just giving it a go ahead, or just be willing to try it out, letting us try it out, it'll be good for me."

4.2.3 Seeking guidance and quality assurance. In some cases, the inherent uncertainty associated with the use of GenAI – for example, when a use case is novel or niche, or when users are lacking in expertise – seeking guidance and clarity from others can also be important for verifying and ensuring the quality of GenAI outputs. P18 explained how she and her supervisor would collaborate over her use of GenAI, stating: "My supervisor is not inclined to AI, and I use AI on most of my work activities, so I usually double check with him [...] and the knowledge or the aspect that I'm using AI on. So that's just to make sure it's all correct, and he's on tow with it and that we understand each other". Similarly, P20, an accountant, shared that: "Despite getting an answer from AI, you still have to confirm, at times from our senior provider, at times from our colleagues [...] for the reliability of our data."

4.3 Socio-technical Competency: Removing cues indicating GenAI use

GenAI outputs were often perceived to be different to human outputs. Participants described different cues which they felt indicated evidence of GenAI use. For some, the cues were identifiable in specific patterns of behaviour and word choices. P21 suggests that GenAI tools display different 'habits' to human developers when coding: "When AI writes the programmes, especially ChatGPT, it adds a lot of comments which normally we developers don't do." For others, evidence of GenAI-generated outputs is detected through a perceived lack of contextual understanding or authenticity. For example, P2, whose work involves writing bids for formal tender opportunities, describes how she can tell when her colleagues have "popped a question in [ChatGPT], it's regurgitated out a piece of writing and then they've sent it over", because the GenAI tool "doesn't have that ability to, (1) read between the lines; and (2) just really understand what points to emphasise." These characteristics were often associated with lower overall quality in work. P9 described: "I've seen content where I feel like, I can tell it's AI generated [...] It

wasn't written by a professional because it's just a little bit average, if you like. It's just not as clever as it should be."

Developing the competency to remove both types of cues of GenAI use is therefore essential for users. We identified two approaches towards cue removal which we differentiate based on participants' motivations. The first is 'concealment', which involves an intentional act of hiding GenAI usage, motivated by social stigma. The second is 'critique', where cue removal is primarily motivated by concerns regarding the quality of the work output, and often driven by a negative assessment of GenAI's capacity to meet the needs of the work. These approaches do not appear to be mutually exclusive, but we highlight the difference in motivation between the two as each can achieve different goals for users, and have different implications for users' willingness to disclose their GenAI use. Below, we highlight how 'critique' can serve as an act of quality control and a behaviour to signal domain expertise and establish role effectiveness, while 'concealment' is predominantly a reaction to perceived social stigma and external judgement.

4.3.1 Critique as quality control over GenAI outputs. Participants who felt they were able to identify and critique the limitations of GenAI-generated outputs often viewed themselves as maintaining important quality control over these outputs. For P2, this behaviour extended to her colleagues' work. By identifying or even gently calling out colleagues for their uncritical usage of GenAI outputs, she could "make the point that [she] know[s] that they've used it [...] to know that [she] knows they haven't done that piece of work [...] and] also to make the point that [she] can pick up on AI quite easily." These interactions could then be used as examples to highlight why it was important to reflect on appropriate reliance on GenAI: "When I talk to people about it at work [...] my point is, AI is not the solution. We can utilise it, but it's not going to replace anything we're doing because it's just not at that level yet." For others, where authenticity and uniqueness of the content generated is an important aspect of their work, incorporating their own efforts and perspective is an important strategy for mitigating the negative implications of 'generic' GenAI outputs. P25, who generates creative content for marketing, explained: "I just don't feel confident enough, just giving whatever [ChatGPT] gives me to someone else. I always want to make sure that I've looked at it with my own eyes and put my own personal stamp on it, where I feel necessary or applicable."

4.3.2 Expert identity is validated through GenAI critique. As GenAI outputs can be viewed as low quality, the opportunity to critique and correct its outputs was also viewed by some participants as opportunity to signal domain expertise to others (see P2 in Section 4.3.1) or to oneself. P19, discussing challenges with debugging when building software using GenAI tools, remarked how it is much easier "if you have expert knowledge of the particular industry or particular problem that you're trying to solve [...] The AI is a bit confused, [so] I'll jump in and I'll fix it." Identifying limitations in GenAI's capacity to address context-specific work can also help to validate individuals' expertise in relation to GenAI. As P25 explained: "I can't say I'm the expert in the industry or the expert in marketing, definitely not. But I have a pretty good understanding of how this company works and what people's expectations are [...] ChatGPT, it just probably doesn't have that." P13, a lecturer who uses ChatGPT in various writing tasks, described how he would use the guidelines for detecting

GenAI use among his students to monitor his own writing and maintain his personal style: *"There are indicator words, as we call them, that students have been using generative AI, so they're words like 'leverage', 'delves', 'underscores', 'deep dives' etc. So when I write, I make sure that those words are not in. [...] So I use [the guidelines], as much as possible, to give my own voice to the information."* He later remarks that: *"I know people have strategies they use on AI that works really well. If you're very clever, no one actually knows you've used it."*

In contrast, lack of domain expertise is primarily associated with over-reliance, or lack of critique, towards GenAI outputs. P4, who worked as a social media marketing manager, described his approach towards training younger professionals as the following: *"If you're new, coming up, I'll say: 'OK, don't rely on it too much or use it too much. You're trying to develop your own creative skills', [...] but if you've been in this work for a while, of course you can use it."* On the other hand, P7, a data analyst, assumed that only he and his colleagues would be using GenAI tools on account of their relatively more junior position, and described his surprise when he discovered that his boss also used GenAI in his work: *"I mean, he's the boss. So I thought he just wouldn't use AI, you know, like he would just come up with all these high-level reports and the likes from his head."*

4.3.3 Concealment as a strategy for managing social stigma. Users were motivated to conceal cues of GenAI usage due to social stigma and potential negative judgements from those around them. P25 shared her speculation of her colleagues' beliefs: *"There might still be a lot of concern about... is [GenAI] accurate? Is it ethical? [...] I've not come across anyone who actively expresses these views, but you never know."* For several participants, discussions of GenAI usage was seen to be potentially controversial and likely to incur negative judgement on themselves or their work, and was thus thought to be safer to avoid. P17 described her dynamics with older colleagues: *"I think with the older generation [...] they look down on generative AI [...] they would say, 'why don't you look at books and do it manually?' [...] They're quite set in their ways, so it's just best not to agitate them by saying that we use a faster, quicker, more efficient method."*

4.3.4 Removing cues of GenAI use negatively impacts access to shared knowledge. Though the ability to critique and conceal GenAI use allowed individuals to maintain quality over GenAI outputs and validate their own expertise, it also had a negative impact on others' access to shared knowledge, and their overall awareness of whether or not GenAI is used. Part of the challenge lies in the fact that users differentiate GenAI usage based on familiarity with colleagues' capabilities, styles, and work outputs, as P4 explained: *"If it's someone I know, I know this person's capability, so when they say 'This is my own work', you know [if] they used [GenAI]".* He acknowledged, however, that this task of differentiation was *"not easy, because [GenAI] is so advanced."* Concerns are exacerbated by environments where GenAI use is not formally acknowledged or established, leading to significant uncertainty and a lack of transparency, even if usage is informally condoned. P10 described: *"I wouldn't want to tell you I use AI for my day-to-day job while we still have these governance and policies and controls in drafts [...] I feel everyone uses AI in one way or another, but no one is saying it to each other because it's not policy enforced yet."* In P25's case, it was

only through direct observation of colleagues using GenAI tools on their screens that indicates to her that these tools are used in other areas of their work.

Breaching the subject of others' GenAI use required significant social sensitivity, while disclosing one's own usage required ignoring potential feelings of embarrassment. P2, who would openly question her colleagues for undisclosed GenAI use, explained how careful she is with her tone: *"When other people will do the same job as me [...] and I look at their responses. I'll be like 'guys, you've just popped this into AI, haven't you?' Because I find it really obvious. [...] the way I approach it, I'll just be like, jokingly, we just sort of have like a laugh about it."* On the other hand, P4, when describing his feelings when sharing his use of ChatGPT with others, explained: *"Where I'm working, I have my device connected to a big screen [...] I don't hide to use ChatGPT, right? I don't hide it. I'm not ashamed of using it."*

However, in other cases, users can refrain from disclosing their experiences with GenAI due to a perceived lack of interest from those around them. This may be shaped by a strong cultural focus on outcomes over processes, which can also stifle discussions around the use of GenAI. P21, a software engineer, summarised his experiences as follows: *"Just get the job done. However you did it, nobody cares. So [...] don't make us talk about it. [...] Maybe it's just indifference, or not the focus for the moment. It's not yet a big deal for us to start adopting it in some ways."* P6, a data analyst whose work was primarily used to support clinicians, therapists, and other admin staff, similarly expressed: *"They [her colleagues] just want the answer [...] They wouldn't be interested in how I did it. I'm ever so excited if people ever are, but they never are."* For some individuals, such norms can make even informal conversations around tool use appear irrelevant, or even socially inappropriate. P15, a professor in digital business, explained that, despite experimenting with GenAI for research analysis, and even developing a course to support PhD students in using GenAI for research, outside of these work contexts, informal conversations about GenAI was *"not suitable. It's not interesting. I don't go into a coffee morning and start talking with people about generative AI. People'd think 'are you crazy?'"* She likened the importance of using GenAI to existing software, stating: *"More generally, I don't talk about [GenAI], because I wouldn't talk about how I use Google or Excel or things like that."* P12, a software engineer who uses GenAI extensively when coding, echoed this sentiment: *"You wouldn't consult your colleague for every line of code that you wrote on your keyboard, for every key press, and I think AI sits at that level."*

The desire to maintain a positive social image can thus negatively impact disclosure and knowledge sharing essential for developing users' technical skills. One key example exemplifying this issue is when P14 described how a colleague asked him to solve a problem she had encountered with Excel formulas. Despite referring to an GenAI tool himself to resolve the bug, P14 chose not to disclose his use of GenAI to his colleague, stating: *"She thought I was a genius. [Asked if he told colleague about using AI] No, no, I did not [tell her]. I think because we didn't talk about how did you get it done. [...] I think that's putting me ahead. Kind of like, they can always come to me [...] like, 'oh, don't worry if you have issues, just walk up to [P14]. He's going to help you out."*

4.4 Strategy for Learning: Learning through Hands-On Experience

4.4.1 GenAI competence as implicit knowledge. Several participants subscribed to the idea that the skills required to support GenAI use is implicit and challenging to articulate. P9 explained: *"It's like, knowing how literal to be, how short and long or accurate your prompt should be, being able to spot when it's not quite right[...] you develop a sense of, 'there's a chance I could have made a mistake in this long block of complex text [prompt]'"* As a result, some participants, such as P15, suggest that GenAI literacy can only be developed through hands-on experience: *"I think it's not something you can read about, you need to test to use them and to learn how to use them to understand the strengths and weaknesses."*

4.4.2 Contextual evaluation of GenAI tools. Participants felt it was important to identify the contextual appropriateness of different GenAI tools and outputs. The extent to which GenAI tools can be appropriately relied on, especially in creative domains, can be highly subjective and contextual, and thus users must individually evaluate and compare GenAI outputs with their own domain knowledge to ascertain which tools are most contextually appropriate or useful. For example, P9, who compared the outputs of different GenAI models for copywriting, reflected on the subtle differences across these tools: *"So we used Copilot, but we also used ChatGPT and [...] Gemini [...], they were all kind of fairly similar, but [...] one might have been better at character counts. One might have been better at [...] creativity and doing funny stuff."* P4 similarly described how he might directly test and compare across different models in order to draw conclusions about what was more appropriate for his context: *"Let me explore with Bard, let me explore with Perplexity. Because I think the algorithms are different [across different GenAI tools]. They kind of do the same thing, [but] like little differences [...]"* Because the ability to evaluate GenAI outputs is closely related to how one perceives one's own domain expertise (as seen in Section 4.3.2), hands-on interaction, experimentation, and comparison are viewed as important for developing both domain knowledge and GenAI competencies simultaneously.

4.4.3 Understanding of technical limitations . GenAI models were often inconsistent in processing prompts and suffered from various limitations. As P2 outlined: *"I've learnt that you have to word things in a certain way, and like, if you put a couple instructions in one sentence, then you might find that the attention is paid to, like, the first part and not the second."* Such inconsistencies meant that errors or hallucinations were unexpected, and were often presented plausibly so that it is difficult to users to detect, requiring caution and direct supervision over model outputs. P6, who used ChatGPT to help her write formulas for spreadsheet calculations, shared: *"When I first started using [ChatGPT], I think 'God, that's amazing. That came back really quickly and I'm sure that's right,' and I'd do it and it wouldn't be right."*

4.4.4 Development of metacognitive understanding. Finally, users had to develop their own metacognitive understanding of how to prompt and articulate their own needs sufficiently to the GenAI model. P9 shared how interactions with tools such as ChatGPT and Copilot helped him to recognise what specifically needed to be articulated. *"I thought that'd be obvious, but now I realise I need*

to state everything that I also don't want [...] I need to also state, what I don't want to happen as a kind of spin off result." For other participants, direct experience and interactions helped them to recognise the importance of having a clearly defined goal prior to interacting with the GenAI tool. P6, who was initially uncritical of ChatGPT's outputs but whose thinking had evolved over time, stated: *"So I'd have to think about what I want it to look like and what I want it to show me.[...] I know enough to question what it's telling me. I can read [ChatGPT's answer] and think, 'well, no, that's not going to work' [...] So then I put another prompt in and it goes backwards and forwards a little bit until we get to where I want it to be."*

5 Discussion

The goal of our study was to extend existing interpretations of GenAI literacy by providing detailed insight into the beliefs, learning strategies, and in particular extending understanding of how users interpret socio-technical GenAI competencies within a sample of 19 knowledge workers. We found that, while participants valued learning from their colleagues' experiences in order to expand their own awareness and knowledge, the ability to remove cues indicating their own GenAI usage was also important. This activity inevitably reduces others' ability to detect GenAI usage, and therefore to develop awareness and access to others' experiences. This promoted a lack of transparency around GenAI use which was exacerbated in organisations with unclear policies. However, we also found that the ability to remove obvious GenAI 'tells' was not always motivated by stigma and fear of social judgement (though this was an important influence), but could also be interpreted as a positive indication of users' domain expertise. These findings together surface a tension between individual identity work and collective learning that has meaningful implications for how we design for and support GenAI literacy in practice.

Our study contributes and extends existing GenAI literacy frameworks by encouraging researchers to reflect more deeply on how competencies for interacting with GenAI manifest for users in context. In the following section, we reflect on and explore the potential implications of our findings.

5.1 Contrasts with Prior GenAI Literacy Frameworks

Our findings challenge current interpretations about the competency needed to detect GenAI outputs. The most recent GenAI literacy framework proposed by Annapureddy et al. [8] positions the ability to identify cues of GenAI-generated outputs as a way to shield individuals from being misled by bias, errors, and misinformation perpetrated by other GenAI users. However, our findings in Section 4.3 demonstrate that users apply their knowledge of GenAI cues to remove evidence of GenAI usage in their own work as well. While previous literature have predominantly focused on AI hiding behaviours as an act of shame [98, 137], and indeed, we also observed social stigma as an important motivator for concealing GenAI use [90, 98, 101], our study highlights how the ability to remove cues of GenAI use can both serve to improve the perceived quality of work, and help users to validate their own domain expertise over those who 'fail' to hide their use of GenAI tools. Indeed,

GenAI hiding could even be applied to accumulate personal reputational credit for participants, as we saw with P14, who utilised GenAI to solve a colleague's problem, but took personal credit for developing the solution.

This suggests a potential need for researchers themselves to re-frame the behaviours around what is currently perceived as GenAI hiding or GenAI shaming, to consider how such behaviours may in fact benefit other user needs. This mirrors similar research into GenAI non-use, where, rather than viewing non-use as an unambiguously undesirable rejection of a 'useful' tool, researchers have instead highlighted how non-use can instead reflect strategic acts of identity protection on behalf of the user, to preserve essential skills and knowledge rather than delegating to the user [20, 126]. Indeed, the term 'workslop' has also been applied to describe 'obviously' generated GenAI content which is shared without consideration or tailoring for recipients' needs, and so, while we highlight the negative impact of GenAI cue removal on transparency for the purpose of learning, we also recognise that this behaviour can mitigate many of the negative impacts of otherwise producing 'workslop', which can damage collaboration, social relations, as well as productivity [76].

While maintaining awareness and understanding of possible GenAI applications and interactions is a key competency according to our participants, and also well-emphasised in existing GenAI literacy frameworks [8, 28, 86], our study highlights how the execution of different competencies can impact each other. Namely, we highlight how hiding of cues indicating GenAI usage is viewed as an important competency for quality control and affirming users' domain expertise, but can impact wider engagement and visibility of practices and others' ability to learn from others' experiences. Taken together, these findings suggest that competencies around GenAI use cannot be understood in isolation, and that recommendations to support GenAI literacy in the workplace would benefit from better understanding how employees experience and address GenAI usage in situated contexts [63].

5.2 The Risks of Non-Disclosure around GenAI Tools

In our study, participants described concerns with disclosing opinions and experiences of using GenAI to those who may judge them negatively for it, as well as navigating potentially negative social conflicts with other colleagues during discussions, even when personal experience and observations told them that GenAI use was prevalent. More recent works echo these findings, showing that self-reports of GenAI usage were up to 40% lower compared to reports of others' use of GenAI, reflecting a significant social desirability bias when it comes to disclosure around GenAI use [62].

As a practice, GenAI hiding can have important consequences for the long-term sustainability of learning through knowledge sharing and supporting norms of usage. While ultimately an individual decision, hiding behaviours and overall lack of disclosure around GenAI usage can contribute to a collective sensemaking process which normalises 'organisational silences' – an absence of discourse, sharing of perspectives and voicing of concerns within an organisation – on GenAI use [73]. The manifestation of such organisational silences can have significantly negative implications

for organisations' ability to develop an accurate understanding of GenAI usage, prevalence, to detect errors, and thus to engage in decision-making and policy-making which can maximise GenAI effectiveness in organisations. Such patterns have previously been observed in more established software. For example, research in spreadsheet learning highlight that individual perceptions around the social acceptability and reputational gains of discussing a software can have negative implications for their willingness to engage in sharing practices [127, 128]. In the spreadsheet context, this lack of communication can lead to significant technical knowledge and skill loss within organisations which pose practical problems when users engage with the artefacts produced through spreadsheet knowledge [108]. In the GenAI context, however, a lack of discourse and GenAI hiding means individuals' knowledge becomes inherently private and invisible to the organisation [129]. This will pose complications for organisations, who may not be able to identify when and where knowledge and productivity is lost, or even what knowledge is lost, nor be able to necessarily replicate such knowledge and the positive outcomes associated with it.

While other AI literacy frameworks have highlighted the broader need for understanding the use of GenAI in social contexts [8, 86], by focusing on knowledge workers' individual perspectives, we highlight how users navigate the (lack of) organisational structures and clear social cues around GenAI use by adopting concealment practices. By developing a deeper awareness of individuals' priorities and experiences, we can support educators, researchers, and management to implement structures and interventions which could better address these challenges.

5.3 The Challenges of Understanding GenAI's Potential

5.3.1 Lack of discoverability. Despite the perceived universality and versatility of GenAI tools in a wide variety of contexts and tasks [8, 100], translating this potential into practice remains challenging for knowledge workers. The experiences of our participants in Section 4.2 showed that, while maintaining awareness of possible applications is seen as an essential technological competency, the discoverability of capabilities beyond users' current usage patterns remains limited without external support [67]. Supporting user discoverability is a well-established problem in the domain of software learning, particularly for complex, feature-rich software [67, 70]. Users can face significant difficulties in discovering and exploring further features and practices beyond those they are used to applying [24, 67, 70]. We argue that there are parallels and lessons which could be drawn from the current literature to help understand how this issue manifests in the context of GenAI tools.

One explanation for what we observed could be that our participants' experiences uniquely reflect the issues faced by users transitioning from traditional, graphic user interface and feature-based software applications, to tools with generative and far more powerful capabilities. A software application is traditionally developed with pre-defined purposes, functionalities, and workflows by the developers [43, 77], which may or may not fully meet the needs of the users. Knowledge workers typically engage in high levels of application-switching in order to address the needs of a

particular task, particularly if applications do not provide the functionalities they need to access [47]. Challenges with discovering novel GenAI use may therefore be an issue of users navigating with these previous constraints, and not recognising the potential for GenAI's capabilities to be extended beyond the existing modalities, tasks, and domains that they currently use it for. For example, work by Renom [92] highlights that users' abilities to re-purpose and craft alternative practices within a tool can often be limited by their biases and their own prior knowledge of other digital environments, as well as interface cues. It may therefore have been challenging for our participants to learn to recognise that GenAI tools do not share the constraints common to other software. In such cases, improving education and understanding of the underlying technological model lowering barriers to exploration could help to mitigate these potential fallacies in understanding.

5.3.2 Challenges to continuous learning around GenAI. Beyond discoverability, exploring GenAI's capabilities can be challenging because it requires continuous, long-term effort. The possible applications of GenAI tools are both extensive and rapidly evolving. Users must therefore contend with updates to existing models can lead to differences and inconsistencies in outputs and task suitability, as well as make sense of multiple competing GenAI tools and increasingly specialised applications of such tools [9, 56]. What's more, in workplace contexts, where there are often time pressures and competing work demands, it can also be difficult for users to maintain learning on top of their existing workloads [59].

Knowledge sharing between peers and colleagues is therefore especially important as a learning process that enables users to navigate resource constraints [59, 119]. In line with theories of communities of practice (CoP) [123], learning can be understood as a socially situated process in which groups of individuals, united by a shared domain, develop and refine practices through sustained participation and interaction. Consistent with prior work [63], our findings show that much of the learning, exploration, and understanding of GenAI practices is socially negotiated. Because there are so many possibilities for interacting with, and adapting GenAI tools to existing or novel workflows, users benefit from the experiences shared by others to help them make sense of these possibilities. However, our findings also extend CoP theories. While CoP emphasises active participation as central to shaping and reinforcing individuals' membership within a community, we find that, users proactively conceal traces of GenAI usage as a way to reinforce their identities as domain experts. In other words, rather than sustaining interactions to reinforce membership and identity as experienced GenAI users, for some users, competing identity priorities and the desire to emphasise own 'domain expertise' drive behaviours which, long-term, threaten the sustainability of GenAI-focused communities and discourse around GenAI use (see Section 5.2). Indeed, our participants emphasised that much of their awareness of GenAI tools is predominantly driven current popular discourse and heightened media representation, and that such interest is not necessarily maintained once usage becomes normalised [125]. It is therefore important to consider the resources and structures which organisations could implement to support AI literacy not only as a one-time phenomenon, but a continuous process of exploration, maintenance,

and transfer for both individuals and the communities they reside in [38].

5.4 Implications

5.4.1 Organisations. For organisational leaders and managers, we emphasise that the development of an open, social workplace environment is one of the most important conditions for supporting effective GenAI learning. To achieve this, managers should establish effective channels for feedback around GenAI usage, and be open-minded to receiving diverse perspectives from employees [73]. In particular, management should take care to recognise how fear of external judgement dictates behaviours such as GenAI hiding, and be cautious in navigating the power dynamics between employees and managers which can exacerbate these issues [73]. Engaging in discourse around GenAI use in informal contexts, and reducing the distance between managers and employees, can also provide individuals with greater psychological safety that can better encourage social learning and improve awareness across the hierarchy of the organisation [35].

We also highlight that passive or implicit allowance of GenAI use may not necessarily encourage discourse. As discussed in Section 4.3.4, despite observations of GenAI use, individuals may still refrain from engaging in knowledge sharing out of concern of potential negative judgement and repercussions from others [63, 137]. In contexts where there is uncertainty, employees can often opt for risk-averse strategies regarding disclosure which exacerbate 'organisational silences' and willingness to report on concerns or problems [73]. This could be of particular concern when onboarding new employees, as a lack of openness and discussion about the use of GenAI can make it challenging for learning and experience to be circulated. Additionally, in blended work environments, AI acts not only as a tool but as a co-worker and co-author, which amplifies the stakes of disclosure, credibility, and ownership [27].

To address these issues, we suggest that organisations should be explicit and targeted in cultivating communities of practice around GenAI use as a priority [124]. Considering the volume and rapidity of GenAI developments, and the scale of possible GenAI use cases, reward incentives could be established to support the different functions needed to support such continuous learning. For example, organisations could address the issue of maintaining awareness of new developments by providing funding for certain employees to trial GenAI tools, experiment with workflows, and develop specialised solutions. Organisations could also draw explicit attention to the issue of GenAI disclosure, and implement tools and templates to support disclosure from other contexts, such as academia or journalism (e.g. [2, 10, 81, 121]), and tailor them to suit the specific contexts where disclosure is mandated within the organisation or when working with external stakeholders. Beyond simply implementing tools for disclosure, encouraging users to reflect on the appropriateness and design of disclosure tools, can also provide an important opportunity for managers to gain a deeper understanding of how GenAI tools are used by workers, how they believe such information should be communicated, and to collectively negotiate and develop a shared understanding of standards for GenAI disclosure within their community.

5.4.2 Design. Echoing prior work on software learning [51, 59, 63, 127], our study finds that GenAI users benefit from shared experiences and artefacts produced by others. Such knowledge helps users uncover new possibilities for GenAI use and adapt their own practices. However, the key challenge which designers must overcome is to identify the level of abstraction and disclosure users are comfortable with sharing regarding their GenAI usage, under what contexts, and with whom. To address this challenge, we suggest that designers should consider both existing interventions to support the communication of GenAI usage patterns, as well as the implementation of socio-technical systems which can encourage repeated engagement in sharing behaviours.

For example, our participants commonly reflected on the idea that sharing of GenAI-related interaction knowledge can be challenging without first-hand experience, due to the unpredictable nature of GenAI outputs. Existing documentation practices of GenAI users on social media forums reflect this complexity, with users often sharing prompts, partial transcripts, screenshots, or highly detailed descriptions of processes outlining workflows, GenAI tools used and corresponding settings, which can be challenging to parse and interpret in context [63, 114]. We suggest that, initially, designers should focus on readily available traces of GenAI interactions, such as prompts, which could be more readily shared or collaboratively developed between users [39, 63]. However, given the evolving nature of GenAI tools [9], the framing around prompt-sharing should be better communicated, not necessarily for the purpose of replicating outcomes through reuse, but instead in revealing the metacognitive strategies needed to structure and articulate tasks and the interactions needed to reach a desired outcome [29, 115]. To address the socio-technical challenges of stigma around GenAI usage, interventions which foreground positive social incentives may help mitigate concerns around sharing. For example, appreciation systems, widely used in open-source communities and platforms such as Wikipedia [109], allow users to give positive feedback and visible signals of appreciation in response to contributions. Recent systems such as HugReports [50] have demonstrated that, embedding such appreciation systems within a user's development workflow and allowing them to signal appreciation for individual packages and modules created by other developers can help contributors feel recognised for their work, and also to serve as feedback for what is useful and what is not. Extending such approaches to shared GenAI artefacts, such as individuals contributing useful prompts, tool integrations, and workflows may both support knowledge exchange and increase the visibility of social rewards associated with sharing.

We highlight, however, that much of the existing interactions with GenAI use remain private and individual [129], and thus more research is needed to identify where there could be useful overlaps and opportunities for collaborative GenAI usage that could support these forms of interactions and system implementations. Designers could follow the approach adopted by Giannakis et al. [42] to develop a better understanding of the levels of abstraction for sharing and disclosure which users are comfortable with, and what design features are needed to support this method of sharing. This topic remains an open space for exploration which will likely continue to evolve as new patterns for GenAI usage and opportunities for user collaboration emerge.

5.5 Limitations and Future Work

While our work provides important insights into how current adopters of GenAI negotiate with the concept of GenAI literacy in the workplace, there still remain several open questions for future works to explore. First, our work identifies the important role of other users in supporting discoverability of GenAI functionalities and applications to work, but it is not known to what extent which GenAI technologies could be applied to fulfill or replace this role, and what the consequences might be. Many current commercial GenAI tools do not yet incorporate clear, explainable AI principles to comprehensively support user understanding of how they may be used in context. Though explainable AI principles can be embedded into design to support the goal of learning, we suggest that the social influence of others on the practice of GenAI usage is likely to persist regardless. Our findings suggest that, regardless of whether or not people actually learn from their colleagues, the act of sharing one's software knowledge can still be an important act for accumulating social credit and signalling domain expertise, and people's understanding of which GenAI competencies matter (if they matter at all), is still influenced by others' attitudes. Greater long-term research on the implications of applying GenAI tools to support GenAI literacy, and the extent to which such observations may manifest in future generations, will be needed to better support our understanding of how we should engage in technological learning.

A further area for exploration is to consider the influence of individual variables and demographics on how users engage in learning. Though we aimed to capture a diverse sample of participants in our study, we did not explicitly examine the influence of individual factors which could also play a role in how susceptible an individual is to specific external or social influences. For example, prior literature suggest that women possess lower self-efficacy, i.e., confidence in their own technical knowledge [134], and gender differences in motivation and types of knowledge hiding that is also impacted by the perceived social roles of both genders [7]. Future work could consider exploring potential associations between factors such as gender, cultural background, technical expertise, or personal variables such as self-efficacy, and personality variables, as well as organisational culture, and intention to share knowledge.

Finally, the findings of this study is bounded by the characteristics of our sample – namely, participants who were already willing adopters of GenAI technology, including those who may be using GenAI against organisational recommendations. As such, common concerns around GenAI use, including issues of fairness and attribution, which are more prevalent concerns for those reluctant to use GenAI [91], were not identified as strongly as issues of work legitimacy, expertise, and reputation. While our participants reflected on who should use GenAI under different contexts, there was less emphasis on the evaluation and decision-making process surrounding when GenAI should be used, a theme which is often more salient in individuals who engage in GenAI non-use [20], or in contexts where GenAI adoption is enforced [111]. While these perspectives fall outside the scope of this study, incorporating them in future works is important for developing a more comprehensive account of GenAI literacy and user–technology relationships.

6 Conclusion

This paper extends current discussions on GenAI literacy through an empirical exploration into how knowledge workers outside of formal education define key GenAI competencies and strategies for learning. We found that users value broad awareness of GenAI's capabilities due to the complexities of the constantly evolving field and to enhance their own competitiveness. Furthermore, learning from the experiences of others is crucial for maintaining this awareness. However, concealing 'obvious' cues of GenAI usage is also valued as a competency, not only to avoid stigma, but as a way of re-affirming users' own domain expertise. We contribute to existing knowledge on GenAI hiding behaviours by highlighting why users may be motivated to engage in hiding, and that these actions can promote a lack of transparency and critical reflection on GenAI use. To mitigate hiding behaviours, we suggest that designers enhance the visibility of user-generated knowledge when using GenAI tools, and that organisations can better emphasise the benefits of maintaining collaborative learning and open discussions to support responsible, transparent use.

7 AI Disclosure Statement

The authors acknowledge the use of Generative AI tools (ChatGPT 5.2) to support language improvements across the manuscript (to refine phrasing and improve readability), and to support brainstorming and ideation based on author's instructions. All outputs were reviewed and verified by the authors. All substantive ideas, descriptions, summaries, and arguments were developed by the authors, who take full responsibility for the final content. The use of generative AI followed relevant ethical guidance. This disclosure statement was created using DAISY (available at: <https://dryoanaahmetoglu.github.io/daisy-disclosure/>).

Acknowledgments

We are grateful for the funding provided by the Engineering and Physical Sciences Research Council [Grant number: EP/W522077/1] and Microsoft Research through its EMEA PhD Scholarship Programme, and the participants for their time and participation. We also thank the anonymous reviewers for their constructive and insightful feedback which helped improve the clarity and presentation of the paper.

References

- [1] Mark S. Ackerman. 2000. The Intellectual Challenge of CSCW: The Gap Between Social Requirements and Technical Feasibility. *Human-Computer Interaction* 15, 2-3 (Sept. 2000), 179–203. doi:10.1207/S15327051HCI1523_5
- [2] Yoana Ahmetoglu, Marios Constantinides, and Anna Cox. 2026. AI Disclosure with DAISY. doi:10.48550/arXiv.2604.02760 arXiv:2604.02760 [cs]
- [3] Sumeyya Akdilek, Ibrahim Akdilek, and Narissra Maria Punyanunt-Carter. 2024. The Influence of Generative AI on Interpersonal Communication Dynamics. In *Advances in Educational Technologies and Instructional Design*, Sanae Elmoudden and Jason S. Wrench (Eds.). IGI Global, 167–190. doi:10.4018/979-8-3693-0831-8.ch009
- [4] Derek A. Albert and Daniel Smilek. 2023. Comparing Attentional Disengagement between Prolific and MTurk Samples. *Scientific Reports* 13, 1 (Nov. 2023), 20574. doi:10.1038/s41598-023-46048-5
- [5] Ali Alkhatib, Michael S. Bernstein, and Margaret Levi. 2017. Examining Crowd Work and Gig Work Through The Historical Lens of Piecework. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, Denver Colorado USA, 4599–4616. doi:10.1145/3025453.3025974
- [6] Omaira Almatrafi, Aditya Johri, and Hyuna Lee. 2024. A Systematic Review of AI Literacy Conceptualization, Constructs, and Implementation and Assessment Efforts (2019–2023). *Computers and Education Open* 6 (June 2024), 100173. doi:10.1016/j.caeo.2024.100173
- [7] Tatiana Andreeva and Paola Zappa. 2023. Whose Lips Are Sealed? Gender Differences in Knowledge Hiding at Work. *Journal of Occupational and Organizational Psychology* 96, 4 (2023), 828–855. doi:10.1111/joop.12444
- [8] Ravinithesh Annapureddy, Alessandro Fornaroli, and Daniel Gatica-Perez. 2025. Generative AI Literacy: Twelve Defining Competencies. *Digit. Gov. Res. Pract.* 6, 1 (Feb. 2025), 13:1–13:21. doi:10.1145/3685680
- [9] Gagan Bansal, Besmira Nushi, Ece Kamar, Daniel S. Weld, Walter S. Lasecki, and Eric Horvitz. 2019. Updates in Human-AI Teams: Understanding and Addressing the Performance/Compatibility Tradeoff. *Proceedings of the AAAI Conference on Artificial Intelligence* 33, 01 (July 2019), 2429–2437. doi:10.1609/aaai.v33i01.33012429
- [10] Edyta Paulina Bogucka, Marios Constantinides, Sanja Šćepanović, and Daniele Quercia. 2025. Impact Assessment Card: Communicating Risks and Benefits of AI Uses. *Proc. ACM Hum.-Comput. Interact.* 9, 7 (Oct. 2025), CSCW301:1–CSCW301:42. doi:10.1145/3757482
- [11] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, and et al. 2022. On the Opportunities and Risks of Foundation Models. doi:10.48550/arXiv.2108.07258 arXiv:2108.07258 [cs]
- [12] Léonard Boussieux, Jacqueline N. Lane, Miaomiao Zhang, Vladimir Jacimovic, and Karim R. Lakhani. 2024. The Crowdless Future? Generative AI and Creative Problem-Solving. *Organization Science* 35, 5 (Sept. 2024), 1589–1607. doi:10.1287/orsc.2023.18430
- [13] Michelle Brachman, Amina El-Ashry, Casey Dugan, and Werner Geyer. 2024. How Knowledge Workers Use and Want to Use LLMs in an Enterprise Context. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–8. doi:10.1145/3613905.3650841
- [14] Virginia Braun and Victoria Clarke. 2019. Reflecting on Reflexive Thematic Analysis. *Qualitative Research in Sport, Exercise and Health* 11, 4 (Aug. 2019), 589–597. doi:10.1080/2159676X.2019.1628806
- [15] Virginia Braun and Victoria Clarke. 2021. To Saturate or Not to Saturate? Questioning Data Saturation as a Useful Concept for Thematic Analysis and Sample-Size Rationales. *Qualitative Research in Sport, Exercise and Health* 13, 2 (March 2021), 201–216. doi:10.1080/2159676X.2019.1704846
- [16] Kelly Caine. 2016. Local Standards for Sample Size at CHI. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, San Jose California USA, 981–992. doi:10.1145/2858036.2858498
- [17] Huajie Jay Cao, Hee Rin Lee, and Wei Peng. 2025. Empowering Adults with AI Literacy: Using Short Videos to Transform Understanding and Harness Fear for Critical Thinking. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–8. doi:10.1145/3706598.3713254
- [18] Astrid Carolus, Martin J. Koch, Samantha Straka, Marc Erich Latoschik, and Carolin Wienrich. 2023. MAILS - Meta AI Literacy Scale: Development and Testing of an AI Literacy Questionnaire Based on Well-Founded Competency Models and Psychological Change- and Meta-Competencies. *Computers in Human Behavior: Artificial Humans* 1, 2 (Aug. 2023), 100014. doi:10.1016/j.chbah.2023.100014
- [19] Dilek Cetindamar, Kirsty Kitto, Mengjia Wu, Yi Zhang, Babak Abedin, and Simon Knight. 2024. Explicating AI Literacy of Employees at Digital Workplaces. *IEEE Transactions on Engineering Management* 71 (2024), 810–823. doi:10.1109/TEM.2021.3138503
- [20] Inha Cha and Richmond Y. Wong. 2025. Understanding Socio-technical Factors Configuring AI Non-Use in UX Work Practices. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–17. doi:10.1145/3706598.3713140
- [21] George Chalhoub and Advait Sarkar. 2022. "It's Freedom to Put Things Where My Mind Wants": Understanding and Improving the User Experience of Structuring Data in Spreadsheets. In *CHI Conference on Human Factors in Computing Systems*. ACM, New Orleans LA USA, 1–24. doi:10.1145/3491102.3501833
- [22] Thomas K.F. Chiu, Zubair Ahmad, Murod Ismailov, and Ismaila Temitayo Sanusi. 2024. What Are Artificial Intelligence Literacy and Competency? A Comprehensive Framework to Support Them. *Computers and Education Open* 6 (June 2024), 100171. doi:10.1016/j.caeo.2024.100171
- [23] Michael Chui, Eric Hazan, Roger Roberts, Alex Singla, Alexander Sukharevsky, Lareina Yee, and Rodney Zimmel. 2023. *The Economic Potential of Generative AI: The next Productivity Frontier*. Technical Report. McKinsey & Company.
- [24] Andy Cockburn, Carl Gutwin, Joey Scarr, and Sylvain Malacria. 2014. Supporting Novice to Expert Transitions in User Interfaces. *Comput. Surveys* 47, 2 (Nov. 2014), 31:1–31:36. doi:10.1145/2659796
- [25] Kaija Marjukka Collin, Sara Keronen, Soila Lemmetty, Tommi Auvinen, and Elina Riivari. 2021. Self-Organised Structures in the Field of ICT: Challenges for Employees' Workplace Learning. *Journal of Workplace Learning* 33, 2 (March 2021), 95–108. doi:10.1108/JWL-10-2019-0124
- [26] Marios Constantinides and Daniele Quercia. 2025. AI, Jobs, and the Automation Trap: Where Is HCI? In *Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work (CHIWORK '25)*. Association for Computing Machinery, New York, NY, USA, 1–8. doi:10.1145/3729176.3729191

- [27] Marios Constantinides, Himanshu Verma, Shadan Sadeghian, and Abdallah El Ali. 2025. The Future of Work Is Blended, Not Hybrid. In *Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work*. ACM, Amsterdam Netherlands, 1–13. doi:10.1145/3729176.3729202
- [28] Andrew Cox. 2024. Algorithmic Literacy, AI Literacy and Responsible Generative AI Literacy. *Journal of Web Librarianship* 18, 3 (July 2024), 93–110. doi:10.1080/19322909.2024.2395341
- [29] Peter Dalsgaard. 2025. Thinking through Prompting: Cognitive Mediation in Human–AI Interaction. In *Proceedings of the 36th Annual Conference of the European Association of Cognitive Ergonomics (ECCE '25)*. Association for Computing Machinery, New York, NY, USA, 1–6. doi:10.1145/3746175.3747192
- [30] Hai Dang, Lukas Mecke, Florian Lehmann, Sven Goller, and Daniel Buschek. 2022. How to Prompt? Opportunities and Challenges of Zero- and Few-Shot Learning for Human-AI Interaction in Creative Applications of Generative Models. arXiv:2209.01390
- [31] Thomas H. Davenport, D.J. Patil, D.J. Patil, and Michael C. Beers. 1996. Improving Knowledge Work Processes. *Sloan Management Review* 37, 4 (1996), 53–65.
- [32] José Osvaldo De Sordi, Marcia Carvalho De Azevedo, Eliane Maria Pires Giavina Bianchi, and Thiago Carandina. 2021. Defining the Term Knowledge Worker: Toward Improved Ontology and Operationalization. *Knowledge and Process Management* 28, 1 (Jan. 2021), 56–70. doi:10.1002/kpm.1647
- [33] Benjamin D. Douglas, Patrick J. Ewell, and Markus Brauer. 2023. Data Quality in Online Human-Subjects Research: Comparisons between MTurk, Prolific, CloudResearch, Qualtrics, and SONA. *PLOS ONE* 18, 3 (March 2023), e0279720. doi:10.1371/journal.pone.0279720
- [34] Peter F. Drucker. 1993. The Rise of the Knowledge Society. *The Wilson Quarterly* 17, 2 (March 1993), 52–72.
- [35] Amy Edmondson. 1999. Psychological Safety and Learning Behavior in Work Teams. *Administrative Science Quarterly* 44, 2 (June 1999), 350–383. doi:10.2307/2666999
- [36] Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. 2023. GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models. doi:10.48550/arXiv.2303.10130 arXiv:2303.10130 [econ]
- [37] Edward W. Felten, Manav Raj, and Robert Seamans. 2023. How Will Language Models like ChatGPT Affect Occupations and Industries? doi:10.2139/ssrn.4375268
- [38] Dana Feng, Bhada Yun, and April Yi Wang. 2026. From Junior to Senior: Allocating Agency and Navigating Professional Growth in Agentic AI-Mediated Software Engineering. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*. Association for Computing Machinery, New York, NY, USA, 1–24. doi:10.1145/3772318.3791642
- [39] Li Feng, Ryan Yen, Yuzhe You, Mingming Fan, Jian Zhao, and Zhicong Lu. 2024. CoPrompt: Supporting Prompt Sharing and Referring in Collaborative Natural Language Programming. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–21. doi:10.1145/3613904.3642212
- [40] Paulo Freire and Donald P. Macedo. 2014. *Pedagogy of the Oppressed: 30th Anniversary Edition* (30th anniversary edition ed.). Bloomsbury Publishing, New York.
- [41] Mingmeng Geng and Roberto Trotta. 2025. Human-LLM Coevolution: Evidence from Academic Writing. doi:10.48550/arXiv.2502.09606 arXiv:2502.09606 [cs]
- [42] Emmanouil Giannakis, Jessalyn Alvina, Andrea Bunt, Parmit Chilana, and Joanna McGrenere. 2022. Promoting Feature Awareness by Leveraging Collaborators' Usage Habits in Collaborative Editors. In *Graphics Interface*.
- [43] Bengt Goransson, Mats Lind, Else Pettersson, Bengt Sandblad, and Patrik Schwalbe. 1986. The Interface Is Often Not the Problem. In *Proceedings of the SIGCHI/GI Conference on Human Factors in Computing Systems and Graphics Interface*. ACM, Toronto Ontario Canada, 133–136. doi:10.1145/29933.30872
- [44] Stuart Hall. 1997. Representation: Cultural Representations and Signifying Practices. In *The Work of Representation*, Stuart Hall (Ed.). Sage, London, 13–74.
- [45] Jan Hatzius, Joseph Briggs, Davesh Kodnani, and Giovanni Pierdomenico. 2023. *The Potentially Large Effects of Artificial Intelligence on Economic Growth*. Technical Report. Goldman Sachs.
- [46] Teresa Heyder and Oliver Posegga. 2021. Extending the Foundations of AI Literacy. *ICIS 2021 Proceedings* (Dec. 2021).
- [47] Amir Jahanlou, Jo Vermeulen, Tovi Grossman, Parmit K Chilana, George Fitzmaurice, and Justin Matejka. 2023. Task-Centric Application Switching: How and Why Knowledge Workers Switch Software Applications for a Single Task. *Graphic Interface* 2023 (2023).
- [48] Huiying Jin and Masato Uchida. 2024. A Three-Year Analysis of Human Preferences in Delegating Tasks to AI. In *Artificial Intelligence in HCI*, Helmut Degen and Stavroula Ntoa (Eds.). Vol. 14734. Springer Nature Switzerland, Cham, 48–66. doi:10.1007/978-3-031-60606-9_4
- [49] Magnus Høholt Kaspersen, Line Have Musaeus, Karl-Emil Kjær Bilstrup, Marianne Graves Petersen, Ole Sejer Iversen, Christian Dindler, and Peter Dalsgaard. 2024. From Primary Education to Premium Workforce: Drawing on K-12 Approaches for Developing AI Literacy. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–16. doi:10.1145/3613904.3642607
- [50] Pranav Khadpe, Olivia Xu, Geoff Kaufman, and Chinmay Kulkarni. 2025. Hug Reports: Supporting Expression of Appreciation between Users and Contributors of Open Source Software Packages. *Proc. ACM Hum.-Comput. Interact.* 9, 2 (May 2025), CSCW099:1–CSCW099:32. doi:10.1145/3710997
- [51] Kimia Kiani, Parmit K. Chilana, Andrea Bunt, Tovi Grossman, and George Fitzmaurice. 2020. "I Would Just Ask Someone": Learning Feature-Rich Design Software in the Modern Workplace. In *2020 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE, Dunedin, New Zealand, 1–10. doi:10.1109/VL/HCC50065.2020.9127288
- [52] Nigel King. 1998. Template Analysis. In *Qualitative Methods and Analysis in Organizational Research: A Practical Guide*. Sage Publications Ltd, Thousand Oaks, CA, 118–134.
- [53] Eunhye Grace Ko, Shaini Nanayakkara, and Earl W. Huff. 2025. "We Need to Avail Ourselves of [GenAI] to Enhance Knowledge Distribution": Empowering Older Adults through GenAI Literacy. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–7. doi:10.1145/3706599.3720032
- [54] Charlotte Kobiella, Yarhy Said Flores López, Franz Waltenberger, Fiona Draxler, and Albrecht Schmidt. 2024. "If the Machine Is As Good As Me, Then What Use Am I?" – How the Use of ChatGPT Changes Young Professionals' Perception of Productivity and Accomplishment. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–16. doi:10.1145/3613904.3641964
- [55] Pawel Korzynski, Grzegorz Mazurek, Pamela Krzykowska, and Artur Kurasinski. 2023. Artificial Intelligence Prompt Engineering as a New Digital Competence: Analysis of Generative AI Technologies Such as ChatGPT. *Entrepreneurial Business and Economics Review* 11, 3 (2023), 25–37. doi:10.15678/EBER.2023.110302
- [56] Yasmine Kotturi, Angel Anderson, Glenn Ford, Michael Skirpan, and Jeffrey P. Bigham. 2024. Deconstructing the Veneer of Simplicity: Co-Designing Introductory Generative AI Workshops with Local Entrepreneurs. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–16. doi:10.1145/3613904.3642191
- [57] Hao-Ping (Hank) Lee, Advait Sarkar, Lev Tankelevitch, Ian Drosos, Sean Rintel, Richard Banks, and Nicholas Wilson. 2025. The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–22. doi:10.1145/3706598.3713778
- [58] Seyoung Lee and Gain Park. 2024. Development and Validation of ChatGPT Literacy Scale. *Current Psychology* 43, 21 (June 2024), 18992–19004. doi:10.1007/s12144-024-05723-0
- [59] Soila Lemmetty and Kaija Collin. 2020. Self-Directed Learning as a Practice of Workplace Learning: Interpretative Repertoires of Self-Directed Learning in ICT Work. *Vocations and Learning* 13, 1 (April 2020), 47–70. doi:10.1007/s12186-019-09228-x
- [60] Soila Lemmetty and Kaija Collin. 2020. Throwaway Knowledge, Useful Skills or a Source for Wellbeing? Outlining Sustainability of Workplace Learning Situations. *International Journal of Lifelong Education* 39, 5-6 (Nov. 2020), 478–494. doi:10.1080/02601370.2020.1804004
- [61] Zhuoyan Li, Chen Liang, Jing Peng, and Ming Yin. 2024. How Does the Disclosure of AI Assistance Affect the Perceptions of Writing? doi:10.48550/arXiv.2410.04545 arXiv:2410.04545 [cs]
- [62] Yier Ling, Alex Kale, and Alex Imas. 2026. Underreporting of AI Use: The Role of Social Desirability Bias. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. ACM, Barcelona Spain, 1–22. doi:10.1145/3772318.3791073
- [63] Haidan Liu, Poorvi Bhatia, Nicholas Vincent, and Parmit Chilana. 2026. Tracing Everyday AI Literacy Discussions at Scale: How Online Creative Communities Make Sense of Generative AI. doi:10.48550/ARXIV.2603.09055
- [64] Jiqun Liu, Jamshed Karimnazarov, and Ryan W. White. 2025. Trapped by Expectations: Functional Fixedness in LLM-Enabled Chat Search. doi:10.48550/arXiv.2504.02074 arXiv:2504.02074 [cs]
- [65] Xin Liu, Longxin Zhang, and Xiaochong Wei. 2025. Generative Artificial Intelligence Literacy: Scale Development and Its Effect on Job Performance. *Behavioral Sciences* 15, 6 (June 2025), 811. doi:10.3390/bs15060811
- [66] Duri Long and Brian Magerko. 2020. What Is AI Literacy? Competencies and Design Considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–16. doi:10.1145/3313831.3376727
- [67] Eva Mackamul, Géry Casiez, and Sylvain Malacria. 2023. Clarifying and Differentiating Discoverability. *Human-Computer Interaction* 0, 0 (Oct. 2023), 1–26. doi:10.1080/07370024.2024.2364606

- [68] Lina Markauskaite, Rebecca Marrone, Oleksandra Poquet, Simon Knight, Roberto Martinez-Maldonado, Sarah Howard, Jo Tondeur, Maarten De Laat, Simon Buckingham Shum, Dragan Gašević, and George Siemens. 2022. Rethinking the Entwinement between Artificial Intelligence and Human Learning: What Capabilities Do Learners Need for a World with AI? *Computers and Education: Artificial Intelligence* 3 (Jan. 2022), 100056. doi:10.1016/j.caeai.2022.100056
- [69] Hannah Mayer, Lareina Yee, Michael Chui, and Roger Roberts. 2025. *Superagency in the Workplace: Empowering People to Unlock AI's Full Potential*. Technical Report. McKinsey & Company. 1–47 pages.
- [70] Joanna McGrenere and Gale Moore. 2000. Are We All In the Same "Bloat"? In *Proceedings of the Graphics Interface 2000 Conference*. 187–196.
- [71] Susan Michie, Maartje M. van Stralen, and Robert West. 2011. The Behaviour Change Wheel: A New Method for Characterising and Designing Behaviour Change Interventions. *Implementation Science* 6, 1 (April 2011), 42. doi:10.1186/1748-5908-6-42
- [72] Microsoft. 2024. *2024 Work Trend Index Annual Report*. Technical Report. Microsoft. 1–38 pages.
- [73] Elizabeth Wolfe Morrison and Frances J Milliken. 2000. Organizational Silence: A Barrier to Change and Development in a Pluralistic World. *The Academy of Management Review* 25, 4 (Oct. 2000).
- [74] Davy Tsz Kit Ng, Jac Ka Lok Leung, Samuel Kai Wah Chu, and Maggie Shen Qiao. 2021. Conceptualizing AI Literacy: An Exploratory Review. *Computers and Education: Artificial Intelligence* 2 (Jan. 2021), 100041. doi:10.1016/j.caeai.2021.100041
- [75] Trinh Nguyen and Amany Elbanna. 2025. Understanding Human-AI Augmentation in the Workplace: A Review and a Future Research Agenda. *Information Systems Frontiers* (March 2025). doi:10.1007/s10796-025-10591-5
- [76] Kate Niederhoffer, Gabriella Rosen Kellerman, Angela Lee, Alex Liebscher, Kristina Rapuano, and Jeffrey T. Hancock. 2025. AI-Generated "Workslop" Is Destroying Productivity. *Harvard Business Review* (Sept. 2025).
- [77] Midas Nouwens and Clemens Nylandsted Klokmoose. 2018. The Application and Its Consequences for Non-Standard Knowledge Work. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, Montreal QC Canada, 1–12. doi:10.1145/3173574.3173973
- [78] Wanda J. Orlikowski. 1992. The Duality of Technology: Rethinking the Concept of Technology in Organizations. *Organization Science* 3, 3 (1992), 398–427. jstor:2635280
- [79] Wanda J. Orlikowski. 1993. Learning from Notes: Organizational Issues in Groupware Implementation. *The Information Society* 9, 3 (July 1993), 237–250. doi:10.1080/01972243.1993.9960143
- [80] Wanda J. Orlikowski. 2000. Using Technology and Constituting Structures: A Practice Lens for Studying Technology in Organizations. *Organization Science* 11, 4 (2000), 404–428. jstor:2640412
- [81] Faraz Forghan Parast. 2025. AI Usage Fact Label Generator. <https://ailabel.netlify.app/>.
- [82] Eyal Peer, David Rothschild, Andrew Gordon, Zak Evernden, and Ekaterina Damer. 2022. Data Quality of Platforms and Panels for Online Behavioral Research. *Behavior Research Methods* 54, 4 (Aug. 2022), 1643–1662. doi:10.3758/s13428-021-01694-3
- [83] Baolin Peng, Michel Galley, Pengcheng He, Hao Cheng, Yujia Xie, Yu Hu, Qiyuan Huang, Lars Liden, Zhou Yu, Weizhu Chen, and Jianfeng Gao. 2023. Check Your Facts and Try Again: Improving Large Language Models with External Knowledge and Automated Feedback. doi:10.48550/arXiv.2302.12813 arXiv:2302.12813 [cs]
- [84] David Pilgrim. 2014. Some Implications of Critical Realism for Mental Health Research. *Social Theory & Health* 12, 1 (Feb. 2014), 1–21. doi:10.1057/sth.2013.17
- [85] Marc Pinski, Martin Adam, and Alexander Benlian. 2023. AI Knowledge: Improving AI Delegation through Human Enablement. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg Germany, 1–17. doi:10.1145/3544548.3580794
- [86] Marc Pinski and Alexander Benlian. 2024. AI Literacy for Users – A Comprehensive Review and Future Research Directions of Learning Methods, Components, and Effects. *Computers in Human Behavior: Artificial Humans* 2, 1 (Jan. 2024), 100062. doi:10.1016/j.chbah.2024.100062
- [87] Volkmar Pipek, Volker Wulf, and Aditya Johri. 2012. Bridging Artifacts and Actors: Expertise Sharing in Organizational Ecosystems. *Computer Supported Cooperative Work (CSCW)* 21, 2-3 (June 2012), 261–282. doi:10.1007/s10606-011-9138-y
- [88] Diana M. Popa. 2024. Critical Exploratory Investigation of AI Consumption, AI Perception and AI Literacy Requirements. doi:10.2139/ssrn.5070600 social science research network:5070600
- [89] Heily Concepción Portocarrero Ramos, Omer Cruz Caro, Einstein Sánchez Bardales, Lenin Quiñones Huatangari, Jonathan Alberto Campos Trigos, Jorge Luis Maicelo Guevara, and River Chávez Santos. 2025. Artificial Intelligence Skills and Their Impact on the Employability of University Graduates. *Frontiers in Artificial Intelligence* 8 (July 2025). doi:10.3389/frai.2025.1629320
- [90] Irene Rae. 2024. The Effects of Perceived AI Use On Content Perceptions. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (CHI '24). Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3613904.3642076
- [91] Nripendra P. Rana, Rajasshrie Pillai, Brijesh Sivathanu, and Nishtha Malik. 2024. Assessing the Nexus of Generative AI Adoption, Ethical Considerations and Organizational Performance. *Technovation* 135 (July 2024), 103064. doi:10.1016/j.technovation.2024.103064
- [92] Miguel A. Renom, Baptiste Caramiaux, and Michel Beaudouin-Lafon. 2023. Interaction Knowledge: Understanding the 'Mechanics' of Digital Tools. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (CHI '23). Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3544548.3581246
- [93] Jana Retkowsky, Ella Hafermalz, and Marleen Huysman. 2024. Managing a ChatGPT-empowered Workforce: Understanding Its Affordances and Side Effects. *Business Horizons* 67, 5 (Sept. 2024), 511–523. doi:10.1016/j.bushor.2024.04.009
- [94] Fardin Ahsan Sakib, Saadat Hasan Khan, and A. H. M. Rezaul Karim. 2024. Extending the Frontier of ChatGPT: Code Generation and Debugging. In *2024 International Conference on Electrical, Computer and Energy Technologies (ICEET)*. 1–6. doi:10.1109/ICEET61485.2024.10698405
- [95] Advait Sarkar. 2023. Exploring Perspectives on the Impact of Artificial Intelligence on the Creativity of Knowledge Work: Beyond Mechanised Plagiarism and Stochastic Parrots. In *Proceedings of the 2nd Annual Meeting of the Symposium on Human-Computer Interaction for Work*. ACM, Oldenburg Germany, 1–17. doi:10.1145/3596671.3597650
- [96] Advait Sarkar. 2023. Should Computers Be Easy To Use? Questioning the Doctrine of Simplicity in User Interface Design. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems* (CHI EA '23). Association for Computing Machinery, New York, NY, USA, 1–10. doi:10.1145/3544549.3582741
- [97] Advait Sarkar. 2023. Will Code Remain a Relevant User Interface for End-User Programming with Generative AI Models?. In *Proceedings of the 2023 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*. ACM, Cascais Portugal, 153–167. doi:10.1145/3622758.3622882
- [98] Advait Sarkar. 2025. AI Could Have Written This: Birth of a Classist Slur in Knowledge Work. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–12. doi:10.1145/3706599.3716239
- [99] Advait Sarkar and Andrew D Gordon. 2018. How Do People Learn to Use Spreadsheets? (Work in Progress). In *Proceedings of the 29th Annual Conference of the Psychology of Programming Interest Group (PPIG 2018)*.
- [100] Wout Schellaert, Fernando Martinez-Plumed, Karina Vold, John Burden, Pablo A. M. Casares, Bao Sheng Loe, Roi Reichart, Sean Ó hÉigeartaigh, Anna Korhonen, and José Hernández-Orallo. 2023. Your Prompt Is My Command: On Assessing the Human-Centred Generality of Multimodal Models. *Journal of Artificial Intelligence Research* 77 (June 2023), 377–394. doi:10.1613/jair.1.14157
- [101] Oliver Schille and Martin Reimann. 2025. The Transparency Dilemma: How AI Disclosure Erodes Trust. *Organizational Behavior and Human Decision Processes* 188 (May 2025), 104405. doi:10.1016/j.obhdp.2025.104405
- [102] Charlotte A. Shahlaei and Ulrika Lundh Snis. 2022. Conceptualizing Industrial Workplace Learning: An Information Systems Perspective. *Journal of Workplace Learning* 35, 9 (Aug. 2022), 1–21. doi:10.1108/JWL-04-2021-0048
- [103] Yijia Shao, Humishka Zope, Yucheng Jiang, Jiaxin Pei, David Nguyen, Erik Brynjolfsson, and Diyi Yang. 2025. Future of Work with AI Agents: Auditing Automation and Augmentation Potential across the U.S. Workforce. doi:10.48550/arXiv.2506.06576 arXiv:2506.06576 [cs]
- [104] Auste Simkute, Lev Tankelevitch, Viktor Kewenig, Ava Elizabeth Scott, Abigail Sellen, and Sean Rintel. 2024. Ironies of Generative AI: Understanding and Mitigating Productivity Loss in Human-AI Interaction. *International Journal of Human-Computer Interaction* (Oct. 2024), 1–22. doi:10.1080/10447318.2024.2405782
- [105] Alex Singla, Alexander Sukharevsky, Lareina Yee, Michael Chui, and Bryce Hall. 2024. *The State of AI in Early 2024: Gen AI Adoption Spikes and Starts to Generate Value*. Technical Report. McKinsey & Company.
- [106] Alex Singla, Alexander Sukharevsky, Lareina Yee, Michael Chui, and Bryce Hall. 2025. *The State of AI - How Organizations Are Rewiring to Capture Value*. Technical Report. McKinsey & Company. 1–23 pages.
- [107] Alexa Siu and Raymond Fok. 2025. Augmenting Expert Cognition in the Age of Generative AI: Insights from Document-Centric Knowledge Work. doi:10.48550/ARXIV.2503.24334
- [108] Justin Smith, Justin A. Middleton, and Nicholas A. Kraft. 2017. Spreadsheet Practices and Challenges in a Large Multinational Conglomerate. In *2017 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. 155–163. doi:10.1109/VLHCC.2017.8103463
- [109] Emma Spiro, J. Nathan Matias, and Andres Monroy-Hernandez. 2021. Networks of Gratitude: Structures of Thanks and User Expectations in Workplace Appreciation Systems. *Proceedings of the International AAAI Conference on Web and*

- Social Media* 10, 1 (Aug. 2021), 358–367. doi:10.1609/icwsm.v10i1.14737
- [110] Sruti Srinivasa Ragavan, Advait Sarkar, and Andrew D Gordon. 2021. Spreadsheet Comprehension: Guesswork, Giving Up and Going Back to the Author. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–21. doi:10.1145/3411764.3445634
- [111] Ashley Stewart. 2025. Microsoft Pushes Staff to Use Internal AI Tools More, and May Consider This in Reviews. 'Using AI Is No Longer Optional'. *Business Insider* (June 2025).
- [112] Yi Su, Kaiyu Zhang, Qili Wang, and Liangfei Qiu. 2023. Generative AI and Human Knowledge Sharing: Evidence from A Natural Experiment. *SSRN Electronic Journal* (2023). doi:10.2139/ssrn.4628786
- [113] Bartosz Surawski. 2019. Who Is a “Knowledge Worker” – Clarifying the Meaning of the Term through Comparison with Synonymous and Associated Terms. *Management* 23, 1 (June 2019), 105–133. doi:10.2478/manment-2019-0007
- [114] Ningjing Tang, Megan Li, Amy Winecoff, Michael Madaio, Hoda Heidari, and Hong Shen. 2026. Navigating Uncertainties: How GenAI Developers Document Their Models on Open-Source Platforms. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*. Association for Computing Machinery, New York, NY, USA, 1–19. doi:10.1145/3772318.3790910
- [115] 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. arXiv:2312.10893 [cs]
- [116] Kiran Tomlinson, Sonia Jaffe, Will Wang, Scott Counts, and Siddharth Suri. 2025. Working with AI: Measuring the Applicability of Generative AI to Occupations. doi:10.48550/arXiv.2507.07935 arXiv:2507.07935 [cs]
- [117] Stephanie M. Tully, Chiara Longoni, and Gil Appel. 2025. Lower Artificial Intelligence Literacy Predicts Greater AI Receptivity. *Journal of Marketing* 89, 5 (Sept. 2025), 1–20. doi:10.1177/00222429251314491
- [118] Michael B. Twidale. 2005. Over the Shoulder Learning: Supporting Brief Informal Learning. *Computer Supported Cooperative Work (CSCW)* 14, 6 (Dec. 2005), 505–547. doi:10.1007/s10606-005-9007-7
- [119] Päivi Tynjälä. 2008. Perspectives into Learning at the Workplace. *Educational Research Review* 3, 2 (Jan. 2008), 130–154. doi:10.1016/j.edurev.2007.12.001
- [120] Priyan Vaithilingam, Tianyi Zhang, and Elena L. Glassman. 2022. Expectation vs. Experience: Evaluating the Usability of Code Generation Tools Powered by Large Language Models. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems (CHI EA '22)*. Association for Computing Machinery, New York, NY, USA, 1–7. doi:10.1145/3491101.3519665
- [121] Kari D. Weaver. 2024. The Artificial Intelligence Disclosure (AID) Framework: An Introduction. *College & Research Libraries News* 85, 10 (2024). doi:10.5860/crln.85.10.407 arXiv:2408.01904 [cs]
- [122] Patrick Weber, Marc Pinski, and Lorenz Baum. 2023. Toward an Objective Measurement of AI Literacy. *PACIS 2023 Proceedings* (July 2023).
- [123] Etienne Wenger and William Snyder. 2000. Communities of Practice: The Organizational Frontier. *Harvard Business Review* 78, 1 (Jan. 2000), 139–146.
- [124] Etienne Wenger-Trayner, Beverly Wenger-Trayner, Phil Reid, and Claude Bruderslein. 2023. *Communities of Practice Within and Across Organizations: A Guidebook*. Social Learning Lab.
- [125] David Gray Widder and Mar Hicks. 2024. Watching the Generative AI Hype Bubble Deflate. doi:10.48550/arXiv.2408.08778 arXiv:2408.08778 [cs]
- [126] Allison Woodruff, Renee Shelby, Patrick Gage Kelley, Steven Rouso-Schindler, Jamila Smith-Loud, and Lauren Wilcox. 2024. How Knowledge Workers Think Generative AI Will (Not) Transform Their Industries. doi:10.1145/3613904.3642700 arXiv:2310.06778 [cs]
- [127] Qing (Nancy) Xia, Advait Sarkar, Duncan Brumby, and Anna Cox. 2025. "How Do You Even Know That Stuff?": Barriers to Expertise Sharing among Spreadsheet Users. doi:10.48550/arXiv.2506.09216 arXiv:2506.09216 [cs]
- [128] Qing (Nancy) Xia, Advait Sarkar, Duncan P. Brumby, and Anna Cox. 2024. The Paradox of Spreadsheet Self-Efficacy: Social Incentives for Informal Knowledge Sharing in End-User Programming. In *2024 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*, 81–88. doi:10.1109/VL/HCC60511.2024.00019
- [129] Qing Xiao, Qing Hu, Jingjia Xiao, Hancheng Cao, and Hong Shen. 2026. Can GenAI Move from Individual Use to Collaborative Work? Experiences, Challenges, and Opportunities of Coordinating GenAI into Collaborative Newswork. doi:10.1145/3772318.3790984 arXiv:2509.10950 [cs]
- [130] Shixian Xie, John Zimmerman, and Motahhare Eslami. 2025. Exploring What People Need to Know to Be AI Literate: Tailoring for a Diversity of AI Roles and Responsibilities. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–16. doi:10.1145/3706598.3713841
- [131] Yuqin Yang, Ying Zhang, Daner Sun, Wenmeng He, and Yantao Wei. 2025. Navigating the Landscape of AI Literacy Education: Insights from a Decade of Research (2014–2024). *Humanities and Social Sciences Communications* 12, 1 (March 2025), 1–12. doi:10.1057/s41599-025-04583-8
- [132] 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 (CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, 1–8. doi:10.1145/3613905.3650867
- [133] Yumi Yi. 2021. Establishing the Concept of AI Literacy. *Jahr – European Journal of Bioethics* 12, 2 (2021), 353–368.
- [134] Kimberly Michelle Ying, Fernando J. Rodríguez, Alexandra Lauren Dibble, and Kristy Elizabeth Boyer. 2021. Understanding Women's Remote Collaborative Programming Experiences: The Relationship between Dialogue Features and Reported Perceptions. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW3 (Jan. 2021), 1–29. doi:10.1145/3432952
- [135] J.D. Zamfirescu-Pereira, Richmond Y. Wong, Bjoern Hartmann, and Qian Yang. 2023. Why Johnny Can't Prompt: How Non-AI Experts Try (and Fail) to Design LLM Prompts. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–21. doi:10.1145/3544548.3581388
- [136] Emily S. Zhan, Maria D. Molina, Minjin Rheu, and Wei Peng. 2024. What Is There to Fear? Understanding Multi-Dimensional Fear of AI from a Technological Affordance Perspective. *International Journal of Human-Computer Interaction* 40, 22 (Nov. 2024), 7127–7144. doi:10.1080/10447318.2023.2261731
- [137] Zhiping Zhang, Chenxinran Shen, Bingsheng Yao, Dakuo Wang, and Tianshi Li. 2025. Secret Use of Large Language Model (LLM). *Proceedings of the ACM on Human-Computer Interaction* 9, 2 (May 2025), 1–26. doi:10.1145/3711061
- [138] Tao Zhou and Songtao Li. 2024. Understanding User Switch of Information Seeking: From Search Engines to Generative AI. *Journal of Librarianship and Information Science* (April 2024), 09610006241244800. doi:10.1177/09610006241244800
- [139] Araz Zitar, Syed Imran Ali, and Nazrul Islam. 2023. Worker and Workplace Artificial Intelligence (AI) Coexistence: Emerging Themes and Research Agenda. *Technovation* 124 (June 2023), 102747. doi:10.1016/j.technovation.2023.102747