Computer Supported Cooperative Work (CSCW) (2026) 35:16
https://doi.org/10.1007/s10606-026-09549-1

RESEARCH

# Exploring How Knowledge Workers Situate an Organisationally Embedded LLM Agent in their Collaborative Work Practices

Kenton O'Hara*[1], Michael Massimi[2] & Miftah Khan[3]
*[1] *University of Bristol, Bristol, UK (E-mail: kenton.ohara@bristol.ac.uk);* [2] *Slack, San Francisco, CA, USA (E-mail: mike.massimi@salesforce.com);* [3] *Plative, Brooklyn, NY, USA (E-mail: mkhan@plative.com)*

Received: 12 September 2025 / Accepted: 12 August 2026


**Abstract.** Since their introduction, large language models (LLMs), have seen widespread adoption by knowledge workers across a range of industries. A growing body of research is exploring this adoption and examining its implications for the future of work. However, much of this research has focused on individual workers with less attention on understanding how these technologies impact the collaborative practices of knowledge work. Furthermore, studies have primarily explored off-the-shelf LLM systems such as ChatGPT which have limited integration into enterprise systems and workflows. As organisations develop and deploy their own LLM agents, grounded in their specific enterprise systems and workflows, there is an important need to understand how they transform collaborative knowledge work. To address this need we present a qualitative study conducted in an IT services organisation that has introduced a custom LLM agent integrated into its enterprise systems, documents, and workflows. We examine how the employees have integrated the technology into trajectories of everyday practices and workflows and how this shapes collaboration in knowledge work.

**Keywords:** CCS CONCEPTS, Human-Centered Computing, Collaborative and Social Computing, Computing methodologies, Artificial Intelligence, Generative AI, Large Language Models, Knowledge Work, AI Agents, Computer-Supported Cooperative Work, Work Practices, Enterprise AI

## 1 Introduction

Since their introduction several years ago, we have witnessed a large scale and widespread adoption of Large Language Models (LLMs) by knowledge workers across a range of different industries. The evolving applications and implications of these technologies, alongside AI more generally, continue to dominate contemporary discourse about knowledge work futures, as industries seek to understand and adapt to their transformative impact (Butler et al. 2023; Butler et al. 2024;

Published online: 24 September 2026

Chui et al. 2023; Hasker 2025). Contributing to this discourse is a growing body of research that aims to understand how knowledge workers across these industries are adopting or wanting to use these systems in the context of their everyday work (Chui et al. 2023; Nonaka and Toyama 2015; Russo 2024; Sako 2024; Simkute et al. 2024; Sun et al. 2024; Vanukuru et al. 2025; Wagman et al. 2025; Woodruff et al. 2024; Yun et al. 2025). These studies offer important early insights into the value and potential challenges of working with these technologies alongside perceptions of how they may change the future of specific kinds of knowledge work. Much of this work, though, is grounded in the adoption of off-the-shelf LLM tools, such as ChatGPT, and in turn has largely been driven by individual knowledge workers seeking efficiency and productivity gains in their own work. As a result, other facets of LLM adoption and use have been under-explored despite constituting a significant part of how we will need to understand these technologies are currently being, and will continue to be, used in practice.

In the first instance, there is a longstanding tradition within CSCW which moves beyond the perspective of the individual and positions knowledge work as an inherently social, collaborative and situated (Hughes et al. 1992; Nonaka 1994; Nonaka and Takeuchi 1995; Orliowski 2002; Orliowski 2006; Schmidt and Bannon 1992; Seely-Brown and Druid 2000; Sellen and Harper 2003). More specifically, these traditions have done much to highlight how collaborative knowledge work is enacted through various configurations of technologies, infrastructures and artefacts (Orliowski 2002; Orliowski 2006; Star and Strauss 1999; Suchman 1994). This suggests that attending to these situated collaborative practices and their relations with emerging technologies, becomes an important part of how we understand their adoption and their future transformation of practice. More recent work coming from these traditions has extended these arguments to new algorithmic and AI-driven systems, urging closer attention to the emergent situated and collaborative practices that surround their use in contemporary knowledge work (Lindley and Wilkins 2023; Wolf and Blomberg 2019). While there have been efforts to study collaborative experiences with LLMs, these have often typically focused on task-specific investigations, such as collaborative brainstorming or co-creation activities (Gao et al. 2024; He et al. 2024; Muller et al. 2024; Shin et al. 2023; Verheijden and Funk 2023). Although such studies provide useful insights into specific collaboration dynamics in the presence of LLMs, they remain largely detached from everyday situated contexts of collaborative knowledge work. There is therefore a need to build on these initial efforts by examining real world adoption of LLMs within ongoing collaborative knowledge workflows.

Secondly, while individuals have been able to rapidly adopt these tools in support of personal productivity, organisational adoption was initially more cautious and constrained. With general off-the-shelf LLM tools, there were inevitable concerns with data privacy and security, hallucination of inaccurate information

and lack of alignment with respect to organisationally specific knowledge such as SOPs, policies, customer and project information etc. (Lauro et al 2024; Wagman et al. 2025). To mitigate these concerns, organisations increasingly developed bespoke and organisationally grounded LLM agents powered by the same underlying models, but grounded and bounded within the documents, systems and data comprising their specific organisational knowledge bases.

Given their relatively recent introduction, there has been limited opportunity to study these systems in situ and to understand their impacts on the collaborative organisation of knowledge work. A recent study by Wagman et al. (2025) offers some exception here in its exploration of an organisationally bespoke LLM agent in a large science organisation. While offering valuable insights into the use of such systems in classic "copilot" mode, as well as their desire for richer use in science workflow automation, their work does not directly address the situated, collaborative practices of knowledge work, nor how these are shaped through interaction with the LLM.

More recently we have started to see a rapid acceleration in enterprise-level LLM deployment, including large scale rollouts of tools such as Microsoft Co-Pilot, Google Gemini, or Slack's SlackBot. Such tools offer significant developments over previous offerings, in their deeper integration with existing enterprise productivity, communication and collaboration infrastructures. Given this emerging infrastructural embedding of these GenAI tools within organisations, there is potential for an even more dramatic transformation of work and collaboration as knowledge workers adapt their collective practices to these systems over time. With the pace and scale of this institutionalisation, there is an urgent imperative for us to understand how these systems reshape work in practice. Crucially, this entails moving beyond the dominant contemporary narratives of productivity and efficiency towards an examination of how these infrastructurally embedded LLM tools impact the practices of collective sense-making, coordination, awareness, documentation and authorship comprising the negotiated order of knowledge work. These concerns raise important socio-technical questions about how AI-mediated documentation and retrieval practices shape organisational knowing, accountability and responsibility in collaborative work. While there will be significant benefits to be gleaned, these will need to be balanced against associated ethical considerations arising in practice.

## 2 Related Work

To help situate our research, we present an overview of related work that motivates the study and provides key conceptual foundations to inform it. While we acknowledge and are deeply informed by a long-standing tradition of workplace studies within CSCW, we focus here on more recent research on GenAI. We

begin by reviewing studies of GenAI in knowledge work followed by work that examines GenAI in the context of collaboration.

### 2.1 GenAI in Knowledge Work

Since their introduction several years ago, LLMs and their associated capabilities have prompted widespread discourse about their implications in knowledge-intensive professions and industries. As these tools are becoming increasingly embedded in everyday professional tasks, there is a growing body of research offering empirical accounts of the evolving landscape of knowledge work (Chui et al. 2023; Russo 2024; Sako 2024; Simkute et al. 2024; Sun et al. 2024; Vanukuru et al. 2025; Wagman et al. 2025; Woodruff et al. 2024; Yun et al. 2025). Together, these studies provide important early insights into the promises and challenges of deploying GenAI tools in knowledge-intensive domains.

Notable here is the research of Woodruff et al. (2024) who presented a qualitative enquiry into knowledge workers across a range of "knowledge" industries. Focusing on early use of off-the-shelf GenAI tools by individual knowledge workers and their expectations for how these tools would transform their industry, the findings highlighted a desire for the tools to reduce the burden of menial work but with the expectation that humans retain oversight and agency in decision-making and strategy-focused tasks. The work also highlighted concerns among knowledge workers about potential deskilling and dehumanisation of work, though there was a gap between worker perceptions of this impact on them relative to the broader narratives of disruption that are reported in the media. As an early study in the area, the work was necessarily oriented to initial adoption practices of individuals and expectations about resulting industrial transformation rather than organisational-level deployment or integration into organisational knowledge bases, processes and tools. Likewise, because analytic focus centred on individual-level impacts, it didn't really explore how GenAI might reshape collaborative knowledge workplace practices.

In a related study, Brachman et al. (2024) explored adoption and desired future use of LLMs by knowledge workers from a large international technology organisation. They highlighted four key areas of LLM usage including in *document creation/ideation*, *information search and analysis*, *advice seeking,* and *task automation*. Their participants also discussed their anticipated value of LLM-based tools in supporting more role-specific documentation (e.g. "Epics" and "User Stories" for software engineers) as well as greater ability to interface and integrate into their own systems and workflows to better leverage their own data. Because this was early work on the topic, Brachman's research was not set up to offer further insights into emergent everyday practices around these use cases. However, as many of these capabilities are now being realised in organisationally tailored LLM assistants, there is a growing need for empirical case studies that examine their impact on collaborative knowledge work in practice.

Other research has more directly evaluated the impact of LLMs on productivity and quality of knowledge work. For example, Brynjolfsson et al. (2025) investigated the impact of a large language model–powered assistant on the work of customer support agents in an enterprise software company. The research found an overall improvement in productivity but noted that such gains were manifest primarily in low-skill and novice workers (through their faster learning), with a minimal effect on highly skilled staff. Accompanying these productivity gains were improvements in customer sentiment, fewer escalations to managers, and reduced employee attrition. Related findings were noted in a study by Dell'Acqua et al. (2023) which explored how GPT-4 affected productivity and quality outcomes on realistic knowledge work tasks by consultants from a global consultancy organisation. Framing their analysis through the concept of the "*jagged technological frontier*" they suggested that GenAI appears to excel at some tasks while failing at others of seemingly similar difficulty. The research showed how for tasks "*within*" this frontier (e.g., creative ideation), LLM assistance led to significant performance improvements, in terms of output quality, task completion and speed. As per Brynjolfsson et al. (ibid), these benefits were most pronounced among less experienced individuals. However, in tasks "outside" the frontier (e.g., nuanced analytical reasoning), performance was found to significantly degrade. Of significance in this work, then, is an emphasis on task-appropriate use and developing user expertise to navigate this evolving technological frontier.

Similar themes can be found in the work of Simkute et al. (2024) and Lee et al. (2025). Simkute et al's (2024) work highlights a potential productivity paradox of GenAI systems whereby for some tasks and users, they can improve productivity while for others they lead to a reduction. In explaining this "productivity paradox" the authors highlight additional factors to consider such as the shift from *production work* to *evaluation work* (e.g. reviewing GenAI output), challenges of poor workflow fit, and task-complexity polarisation whereby easy tasks are made easier and hard tasks harder. Lee et al. (2025) develop these ideas by focusing on critical thinking in knowledge work and introducing the notion of "*AI stewardship*", emphasising the need for knowledge worker skill development and support tools to guide and monitor GenAI to achieve high quality outputs in their work. Agan, though, much of the analytic focus remained with the individual knowledge worker interacting with the AI tool. While there is acknowledgement of broader workflows, the work is much less explicit about the collaborative trajectories (Strauss 1985, 1988) of knowledge work that span across multiple systems, workers and organisational processes. As we shall discuss through our own work, this omission has important implications for how we both understand and design for such AI stewardship.

Other more recent research (Yun et al. 2025) identifies critical challenges faced by knowledge workers when using generative AI for data navigation and decision-making. In particular, it highlights the difficulties in synthesising large

volumes of unstructured data across platforms, lack of transparency and coordination in collaborative workflows, and a limited capacity to integrate domain knowledge and evolving priorities into AI-driven tools. Through the development of *Yodeai*, the authors explored how AI could potentially assist in structuring, exploring, and prioritising information for product managers. Based on their work, they concluded that AI tools in knowledge work would need to be adaptable to diverse and evolving workflows and personal preferences of the knowledge worker. In addition, they highlight the need for interoperability whereby AI tools need to integrate with existing platform and collaborative contexts in ways that are accountable and source linked. While their tools have not yet undergone any evaluation in real-world contexts, their conclusions resonate strongly with some of our own findings of real-world deployment and use.

## 2.2 GenAI in Collaboration

In terms of understanding GenAI use as a more collaborative concern, much less attention is given to this in the literature. Where there has been consideration of collaborative concerns, this has often been through the lens of the AI as a collaborating colleague in the form of *human-AI collaboration* or *human-AI teaming* (Duan et al. 2024; Duan et al. 2025; Flathmann et al. 2024; Guo et al 2024; Hauptman et al 2022; He et al. 2024; Schmutz et al. 2024; Zhang et al 2023). Much of this research has focused on specific elements of collaboration dynamics between humans and agents such as differences in communication strategies, social influence, trust, interaction and agency. While these dynamics are important, they are not primarily concerned with the everyday situated practices of GenAI use through which collaborative knowledge work is organised.

Other research has examined GenAI as a tool used within human collaboration tasks without positioning the AI as an explicit team member (Shin et al. 2023; He et al. 2024; Gao et al. 2024; Muller et al. 2024. These studies have typically focused on single bounded synchronous tasks and often within the context of design-oriented activities such as brainstorming (Muller et al. 2024) ideation (Shin et al. 2023; He et al. 2024), co-creation (Verheijden and Funk 2023) and qualitative analysis (Gao et al. 2024). While representative of certain knowledge work activities, these tasks are treated as isolated events and not situated within the longitudinal trajectories of interdependent communication, coordination, sense-making and production activities that comprise organisational knowledge work. As Li et al. (2025) note in their systematic review of Human-AI collaboration in Information Systems Research, there is a pressing need for ethnographic and longitudinal accounts of AI use in collaborative work settings.

Building on Strauss' concept of work trajectories (Strauss 1985, 1988), CSCW research has long emphasised the temporal and socio-material unfolding of collaborative work. Central to this work is the notion of articulation work, the work done to keep trajectories on track through the communication and coordination

of interdependent tasks, actors, resources and information. It includes planning, negotiating, scheduling and adaptation to emergent contingencies. Key to early CSCW work, these ideas did much to highlight the emergent, contingent and situated nature of knowledge work and the need for technologies to support this (Schmidt and Bannon 1992; Suchman 1994). Schmidt and Bannon (1992) notably point to the importance of standardisation and structure in organisational processes and artifacts in supporting articulation work along collaboration trajectories. Despite the enduring significance of these concepts there remain significant gaps in how they are applied to the infrastructural embedding of GenAI in contemporary knowledge-based organisations. While the concepts of trajectories and articulation work transformed our understanding of the situated practices of knowledge work, there are currently important gaps in how we connect these ideas to the evolving infrastructural embedding of GenAI tools in knowledge-based organisations. As LLM agents become increasingly implicated in documentation and shared sense-making practices and as they embed within standardised organisational process and artefacts, we would expect a gradual reshaping and redistribution of articulation work. How these tools render aspects of articulation work visible or invisible, enabled or constrained, remain a significant but largely unexplored concern.

More recent research has started to revisit these concepts in the context of emerging GenAI capabilities. Wang et al. (2024), for example, introduced the idea of *Meeting Bridges,* aiming to articulate key principles for linking the articulation work of synchronous meetings with that of asynchronous collaboration. The research highlighted the importance of capturing information from meetings as personal and group archive, for action-oriented reminders, onboarding and inclusion, subsequent sense-making and as starting context for follow-on collaboration. However, they also identified challenges with information capture during meetings, impact on engagement, along with difficulties re-establishing context or finding desired information. The work established key design principles for 'meeting bridge' systems, including incorporation of different linked data types, structure to support information exploration information, customisation information and artifacts, connection to context outside meetings, and the ability to evolve information after meetings. The authors discuss how emerging LLM capabilities could address these challenges and principles. The work provides important conceptual foundations but stops short of developing technologies and evaluating any use in practice. In this respect, there remain important gaps in our understanding of how these tools become infrastructurally embedded in organisational contexts and process and how knowledge workers collectively adapt their practices in use.

In related work, Asthana et al. (2025) further explored these ideas through an LLM-powered meeting recap system as a collaborative artifact for articulation work, documenting discussions and outputs to support future

collaboration and communication. The authors emphasised the importance of multiple recap levels (Highlights and Hierarchical models) that could adapt to different contexts of follow-up work. These recaps are seen as natural artefacts for collaboration, which when seamlessly integrated with enterprise collaboration tools (e.g. Microsoft Teams of Slack) have the potential to transform everyday practices of knowledge work. However, as they acknowledge, there remains a gap in understanding how such LLM-mediated documentation is adopted and appropriated in real everyday collaborative trajectories of knowledge work.

Vanukuru et al. (2025) build on these ideas to examine the temporal nature of knowledge work. Drawing on concepts of *retrospections* (looking back) and *prospections* (looking forward), the research proposed GenAI enabled concepts (*Instant Recaps*, *Adaptive Meeting Handoff* and *Project Browsers*) to support a chain of intentionality and reflection across meetings, where meeting summaries enable reflective retrospection as well as reconnecting to the future with a view to planning and preparing for discussion in upcoming meetings. While these concepts offer compelling design inspiration, there has been limited opportunity to deploy and understand their use in real organisational contexts. In this respect, the evolving situated practices and considerations associated with LLM-mediated prospections and retrospections remain under-examined.

Approaching these issues from a different perspective, Lindley and Wilkins (2023) explored opportunities and considerations for the adoption of machine learning and GenAI in knowledge-based organisations. Their concerns were with the use of AI to generate organisational knowledge bases drawing on the actions and practices comprising everyday knowledge. They examined how AI-infused knowledge bases might enable more effective knowledge work practices. Based on longstanding discussions of organisational knowledge and knowledge work (Nonaka 1994; Nonaka and Takeuchi 1995; Nonaka and Toyama 2015; Orlikowski 2002; Orlikowski 2006; Seely-Brown and Duguid 2000; Sellen and Harper 2003), they frame knowledge work in terms of *actions* rather than *outputs* with *Collaborative Actions* (CAs) and *Knowledge Actions* (KAs) highlighted as potential signals for building ML knowledge bases. These included seeking help, gauging expectations, approvals, planning and strategy work, coordination and articulation work (cf. Schmidt and Bannon 1992; Strauss 1985, 1988; Suchman 1994) required to stay in sync, generate and capture information and transfer information. In acknowledgement of the tacit, embodied and socially enacted nature of organisational knowledge, they argue that machine learning systems could learn both from explicit inputs comprising the organisational knowledge base *and* implicit signals in the socially enacted practices of their use. In this respect, machine learning tools are considered as something to work with *and* as something to work through (Bødker 1989), highlighting the need to examine how such systems reshape the *processes* of

doing knowledge work, in particular in real world organisations. Empirical studies at this level of situated practice, however, remain rare.

Adoption of these tools at a more organisational level was initially slower and more cautious than their individual adoption by knowledge workers owing to concerns of data privacy and security as well as with output specificity. Such concerns are being increasingly addressed through the deployment of enterprise oriented LLMs as well as the development of bespoke LLM agents grounded in organisational knowledge bases. Research into how these tools are used in practice remains limited. An exception here is Wagman et al. (2025) who undertook a study of an organisationally bespoke LLM-powered chatbot, Argo, deployed in the National Science Lab in the US. The research examined perspectives on initial adoption, uses, potential concerns about uses and envisioned future use. The research highlighted a need to work with structured forms of documentation and reporting that were bound to particular processes of scientific practice, with related significance to our own findings. Likewise, they identified an envisioned need to mine insights across large unstructured sets of documents, scientific literature and public data as well as from meeting transcripts, organisational procedures and policies. These envisioned participant needs align closely with key capabilities of the LLM agent in our own study but could not explicitly be examined in Wagman's study where there was less concern with temporal and collaborative features of knowledge work and an analytic orientation to situated organisation of collaborative practices with these LLM agents.

The work of Wolf and Blomberg (2019) offers some further insights here. While this predates the recent introduction of LLMs, it examined the development of a related NLP (natural language processing) infused tool in an IT service provider organisation a setting closely aligned with that of the current paper. Of significance is their analytic orientation, which draws on foundational CSCW approaches to examine the "unfolding and tangled relations between situated work practices and the emergent technologies that we might design to support, augment, or transform such practices." Their analysis presents a rich description of these situated practices and worker accounts of how the developing NLP tool would relates to these. Notable here was the concern that IT architects' envisioned integration with these algorithmic capabilities was collaborative in two ways. Firstly, in terms of human-AI interaction where AI capabilities and outputs would be leveraged and transformed through human interaction. Secondly, that IT solutioning work was seen as a "team sport" that would require algorithmic outputs to be integrated so as to enable their collective interpretation and transformation through the practices of the team.

Taken together, this body of research demonstrates substantial progress in understanding individual-level impacts of LLMs and bounded forms of AI-mediated collaboration activities and dynamics. However, there remains a critical empirical in understanding how LLM-based agents, when infrastructurally

embedded within organisational tools, documentation practices, and project workflows, become implicated in the situated organisation of collaborative work. Existing studies largely abstract away from the longitudinal, trajectory-based nature of knowledge work, offering limited insight into how LLM-mediated contributions are appropriated, negotiated, and aligned across interdependent actors, artefacts, and organisational processes. In particular, there is insufficient attention to how such agents reshape important aspects of articulation work - how coordination, accountability, sense-making, and continuity are accomplished as work unfolds across meetings, communications, documents, and systems. As interaction with enterprise-oriented LLM agents move beyond episodic use towards deeper integration within organisational infrastructures, there is a pressing need for fine-grained, practice-oriented studies that examine how collaborative work is reconfigured in situ.

To address this gap, we present a qualitative study of a professional IT services organisation that has introduced an LLM agent across its various enterprise systems and workflows and that has become integrated into their everyday collaborative work. Specifically, we adopt a practice-oriented perspective to understand the in-situ ways in that LLM-mediated documentation, sense-making, coordination and accountability, are collectively accomplished across interdependent actors, organisational artefacts, processes and trajectories. In doing so we highlight potential benefits of these systems and how these are bound to their infrastructural embedding. Moreover, through the detailed accounts of these emergent situated practices, we highlight potential ethical considerations that arise in the use of these systems and how these are mitigated and managed through the Organisation's cultural orientation. We begin with an overview of the Organisation and the LLM-based agent, Plaito, before presenting the findings and discussing their implications for understanding AI-mediated collaboration in knowledge work.

## 3 Study

### 3.1 The Organisation

Plative is a professional services organisation that develops customised solutions across leading Customer Relations Management (CRM) and Enterprise Resource Planning (ERP) software. They work with clients to understand their industry-specific challenges and pain points and help them transform their systems and processes to address these through CRM and ERP customisation. They work with a variety of clients with specialist expertise in financial services, professional services, nonprofit organisations, technology, media, and consumer sectors. They have a global presence with offices and clients distributed across five continents.

The organisation has approximately 200 employees comprising, consultants, business analysts, developers, sales reps, project and delivery managers, and human resources. While there are office locations, the employees primarily work remotely, embracing a strong culture of online collaboration. The employees therefore rely heavily on *Slack* and *Google Workspace* (Meet, Calendar, Drive, Docs etc.) for their collaboration. For each customer account, there is a dedicated Slack channel for internal team communication and collaboration, and a related external-facing Slack Connect channel that provides a collaborative workspace with the client.

As part of their everyday work and operations, the organisation also makes use of various cloud-based enterprise software tools to manage business operations, collaboration, data management, project management, software development and decision-making These include: *Confluence*—their deep knowledge base covering sales content, delivery content and talent and acquisition content; *Tableau* – a platform for analysing and visualising business intelligence data; *Salesforce* – a CRM tool for managing their customer relationships, sales and marketing; *NetSuite* – an ERP tool for managing core business functions such as finance, HR and inventory; and *Jira* – for project management and tracking tasks/issues in software development workflows.

Such practices reflect a digitally literate workforce that embraces new forms of working and trusting culture among its employees.

### 3.2 "Plaito" LLM Bot

Plaito is a customised AI bot developed by the organisation. It is based on Open AI but tailored and grounded in Plative's knowledge base, enterprise tools and data as well as having broader access to web content (such as LinkedIn profiles). Plaito is available as a standalone tool but also accessible directly from within Slack, where it is integrated in the communication and collaboration channels that underpin the everyday work of the organisation. Employees interact with Plaito via natural language messages and prompts much as they would with other LLM-based tools such as ChatGPT, Llama and Gemini, or if they were directly messaging a colleague in Slack. As with these other tools, the message content can include URLs, links to documents or specific content linked within the enterprise software tools. As well as the natural language prompt input box, Plaito also has a number of "short-cut" buttons (see Fig. 1) including [Ask Plaitive], [Ask Salesforce], [Ask NetSuite], [Ask BI], [Call Prep], [View Portfolio], [Ask Jira], [My Tickets]. These short-cuts provide quick structured access to frequently required information from the various enterprise tools and online sources (Confluence, Salesforce CRM, Netsuite ERP, Tableau BI, LinkedIn and Jira). For example, the 'Call Prep' button provides a quick way to glean information about client participants identified in Google Calendar invites by summarising key insights from their personal and organisational LinkedIn profiles. The

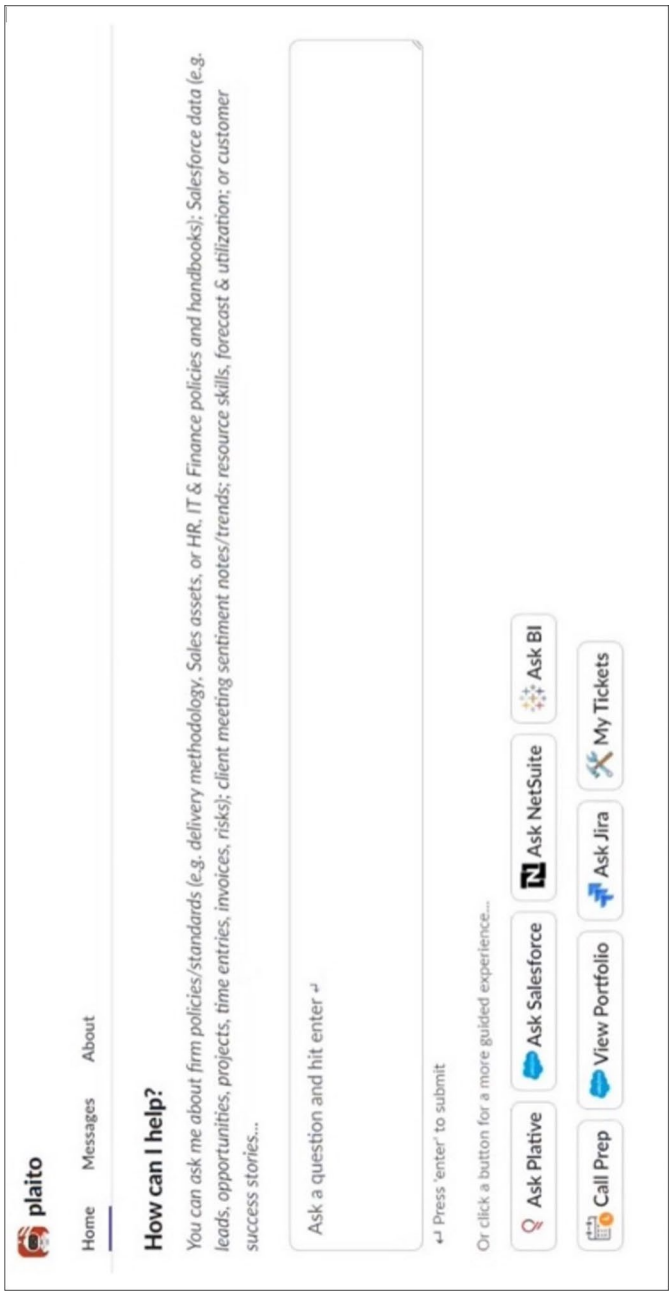


**Fig. 1** Plaito AI bot interface including natural language prompt input and shortcut buttons

"View Portfolio" button provides a listed overview of projects that an individual is working with links to further information about each.

A further key function of Plaito relates to the recording and transcription of meetings. By default, all meetings are recorded and transcribed and then summarised in a standard format by Plaito. The summaries are available from within the Slack channels and Google Calendar events. The meeting transcripts are stored and are accessible to Plaito for subsequent analysis and prompt-based interrogation. Of note here is that the meeting transcripts are subject to a sentiment analysis of the language to captures a sense of the overall feeling of the meeting and ongoing relationship, rating it with a score of positive or negative sentiment. These sentiments ratings are further augmented as a traffic light system (red, amber, green) associated with each project and meeting.

### 3.3 Interviews

To orient the research team, a senior leader in Plative provided us with an introductory overview of the organisation, the nature of its work, and the key stakeholder roles involved in delivering their services. The key stakeholder roles identified were sales leaders, sales team members, delivery leaders, project managers, functional consultants, junior analysts, developers, and HR. An internal communication was sent to various team leads of the functional roles to inform them of the study and to help identify potential representatives of these roles that we could invite to participate in the research. As well trying to recruit across these roles, we were also hoping to identify certain clusters of individuals that worked with each other on projects with a view to representing existing collaborative relationships. Based on these criteria, and recommendations from team leads, we collectively identified a shortlist of 22 potential candidates to approach about study participation, of which 14 agreed to participate (see Table 1).

An initial semi-structured interview (approx. 30 min) was conducted with all fourteen participants. During these interviews, we discussed their role at Plative, their everyday workflows, activities and tasks, the key tools, systems and artefacts they use and create in the course of their work, who they collaborate with and how Plaito more specifically was integrated into the collaborative aspects of their work.

Following this interview, participants were, availability permitting, invited to participate in a second interview two weeks later – nine were available to continue further. The purpose of this second interview was to enquire more deeply about some specific concrete episodes of Plaito use in the context of workplace collaboration. During these two weeks, and in preparation for their interview, participants were asked to choose three to five concrete episodes of Plaito interactions they had undertaken in the course of their everyday work. More specifically, participants were requested to log those examples where interactions with Plaito were embedded in a larger collaborative engagement with one or more

**Table 1** Participant roles and interviews

| Participant | Role | 1st interview | 2nd Interview |
|---|---|---|---|
| P1 | Principal Consultant | Y | Y |
| P2 | Senior Director | Y | Y |
| P3 | Program Manager | Y | N |
| P4 | Senior Project Manager | Y | Y |
| P5 | Delivery Manager | Y | N |
| P6 | Consultant | Y | Y |
| P7 | Business Analyst | Y | Y |
| P8 | Business Development Manager | Y | Y |
| P9 | Senior Director | Y | Y |
| P10 | Sales Rep | Y | Y |
| P11 | Talent Acquisition Manager | Y | Y |
| P12 | People Business Partner | Y | N |
| P13 | People Business Partner | Y | N |
| P14 | Talent Acquisition Manager | Y | N |

of their colleagues or had a material bearing on how they organised their collaboration with clients and colleagues. Participants were further asked to document these episodes with screenshots of their Plaito interactions and any relevant documents/systems that formed part of the larger collaborative engagement in which they were situated. The second interview (approx. 60 min) was structured around these collected episodes, systematically working through each episode in turn. Using their logs as an aide-memoire, participants were asked to recount the details of these episodes and contextualise them within the larger workflows and tasks in which they were embedded. Where appropriate, participants would share their screen to show supporting material, such as the unfolding Plaito interactions, any documents produced, any related system interactions and associated communications with team members and clients. Based on noteworthy features within participant accounts and the shared artefacts, interviewers would interrogate participants to further unpack situated details of the episodes, in particular to help draw out features relating to the collaborative concerns. Over and above the rich descriptions provided, participants were also asked to offer explanatory reasoning of why these situated activities unfolded in specific ways, again with reference to the people they were working with, artefacts used, and broader organisational considerations. Specific attention in these interrogations was given to exploring the details of their iterative prompts to Plaito, and the interactive work involved to create and utilise meaningful outputs that addressed certain challenges or reconfigured the situated organisation of their collaborative activity. These would be contrasted with former ways of doing things to more explicitly understand the nature and value of these work transformations.

All interviews were conducted online using Microsoft Teams or Google Meet. Automated transcripts were produced using inbuilt transcription software, subsequently reviewed, corrected and anonymised in preparation for analysis.

### 3.4 Analysis

The interview transcripts were analysed using a reflexive thematic analysis (RTA) approach (Braun and Clarke 2006) acknowledging an active interpretative role of the researchers as fundamental to the meaning making and analytic account produced. We adopted an abductive approach that was primarily inductive data-driven with additional deductive elements drawing on sensitising concepts from CSCW to develop the interpretation. This allowed us to remain grounded in the experiential perspectives and understandings of our participants while using concepts and theory to help make sense of the patterns in the data. We followed key stages of the RTA process, beginning with data familiarisation, in which we reviewed the original recordings, corrected the transcripts, and followed up with several passes through the transcriptions. During these reviews, tentative interpretive reflections were noted within the transcripts. After this, we performed an additional pass through the transcripts, identifying and labelling salient extracts that related to the situated use of Plaito embedded in a larger collaborative sequence of work with clients and colleagues. Meaningfulness and significance in the data were given extra weight with recurrence of issues across different participants. With these initial codes, attention was primarily on capturing and describing interactional details as reported by participants. Following this, we shifted to a more interpretive phase of the analysis. Codes were clustered together to construct initial high-level patterns in the data. For each of these higher-level patterns, we then iteratively engaged with key concepts and theoretical concerns from CSCW, such as the situated negotiated accomplishment of work, coordination activity, awareness, articulation work, reflexive accountability, and the spatial and temporal ordering of work along trajectories. Through these iterative engagements we were able to derive higher level theoretically informed thematic interpretations of the patterned ways in which collaborative knowledge was reconfigured and accomplished with the Plaito LLM agent.

## 4 Findings

The findings are organised around the key themes emerging from the thematic analysis. We begin with an exploration of how meetings become documents and how this transformation impacts collaborative practices both during meetings, and afterwards along the trajectory of subsequent work. We continue with a look at how these “meetings” become interrogated, transformed and remixed into new knowledge-based materials with the suite of enterprise knowledge tools.

Following this, we unpack how these capabilities create new opportunities to efficiently reconfigure how collaboration with colleagues and clients is enacted. We begin with a look at how meetings are materialised as documents and the influence of this on the organisation of collaboration within the meeting. We then discuss how the subsequent interrogation of the meeting enabled by this creates new possibilities for configuring collaborative work, spatially and temporally. Following this, the findings explore how the materialisation of meetings and agential interrogation shifts the ontology of collaborative authorship and enables the "remixing" of this materialised content into new workflow specific systems and forms. As a counterpoint to "meetings as documents", we further consider how employees can introduce their own workflow and process improvements by generating "documents as meetings". Finally, we discuss new ways in which situational awareness across multiple projects, clients and resources is enabled and the importance of this in the nurturing of ongoing professional client relationships.

### 4.1 Materialising Meetings as Documents

Given the distributed nature of team members and clients, almost all meetings are conducted online via Google Meet. Such an on-line first culture has enabled the organisation to adopt a specific stance towards the recording of meetings. By default, the culture within the organisation is that all meetings are recorded and automatically transcribed with LLM-enabled transcription software. While larger, more traditional, and less digitally homogeneous organisations may encounter resistance to the widespread adoption of such practices, this organisation has embraced them in pursuit of operational and collaborative benefits. As we shall see, however, these practices also introduce new forms of accountability and situational oversight that depend upon careful deployment, trust, and an open organisational culture. To illustrate, consider the role of Plaito. From the meeting transcript, Plaito automatically generates a structured summary document. These summaries are attached to the associated Google Calendar event and stored within the organisation's Google Drive knowledge base. As shown in Fig. 2, the summaries follow a standardised format defined by the including sections such as *Meeting Title*, *Date*, *Time*, *Attendees*, *Overall Summary* and *Key Takeaways*, *Key Decisions*, *Issues/Blockers*, *Action Items/Next Steps*, and any *Additional Context*. Each of these sections is populated by a bulleted list of the key points.

Of note in these accounts is how the Organisation defines what constitutes "organisationally-relevant" information within meetings but also the agency delegated to Plaito in producing these accounts. Because the summaries are attached to Calendar entries, and archived in organisational systems, they have the potential to be treated as "official" versions of events. Any uncontested deferral of responsibility to Plaito for such accounts brings inherent risks, something acknowledged by the Organisation as evidenced in a highlighted disclaimer about

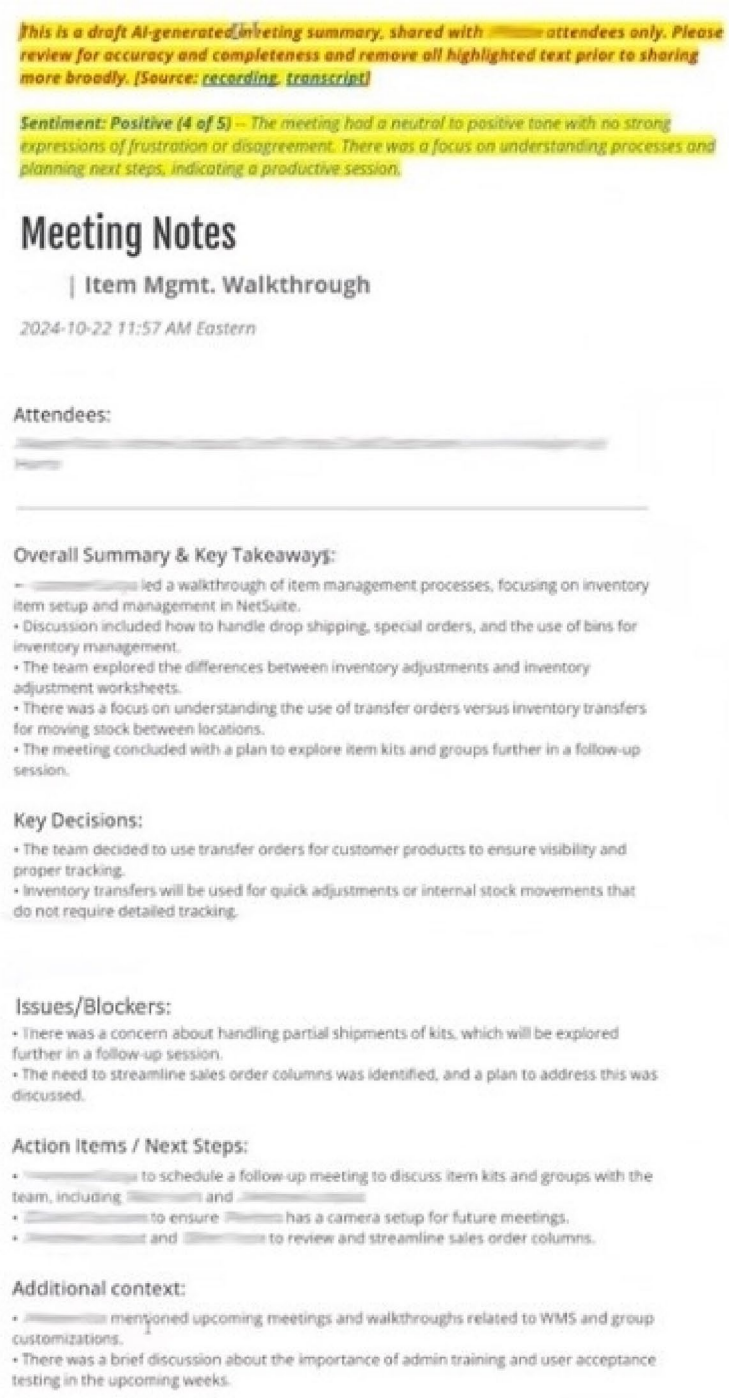

*This is a draft AI-generated meeting summary, shared with [redacted] attendees only. Please review for accuracy and completeness and remove all highlighted text prior to sharing more broadly. [Source: recording, transcript]*

*Sentiment: Positive (4 of 5) – The meeting had a neutral to positive tone with no strong expressions of frustration or disagreement. There was a focus on understanding processes and planning next steps, indicating a productive session.*

# Meeting Notes

| Item Mgmt. Walkthrough

*2024-10-22 11:57 AM Eastern*

## Attendees:

[redacted]

## Overall Summary & Key Takeaways:

– [redacted] led a walkthrough of item management processes, focusing on inventory item setup and management in NetSuite.
• Discussion included how to handle drop shipping, special orders, and the use of bins for inventory management.
• The team explored the differences between inventory adjustments and inventory adjustment worksheets.
• There was a focus on understanding the use of transfer orders versus inventory transfers for moving stock between locations.
• The meeting concluded with a plan to explore item kits and groups further in a follow-up session.

## Key Decisions:

• The team decided to use transfer orders for customer products to ensure visibility and proper tracking.
• Inventory transfers will be used for quick adjustments or internal stock movements that do not require detailed tracking.

## Issues/Blockers:

• There was a concern about handling partial shipments of kits, which will be explored further in a follow-up session.
• The need to streamline sales order columns was identified, and a plan to address this was discussed.

## Action Items / Next Steps:

• [redacted] to schedule a follow-up meeting to discuss item kits and groups with the team, including [redacted] and [redacted]
• [redacted] to ensure [redacted] has a camera setup for future meetings.
• [redacted] and [redacted] to review and streamline sales order columns.

## Additional context:

• [redacted] mentioned upcoming meetings and walkthroughs related to WMS and group customizations.
• There was a brief discussion about the importance of admin training and user acceptance testing in the upcoming weeks.

**Fig. 2** An example of a meeting summary generated by Plaito

AI-generated summaries that should not be distributed beyond the immediate attendee list or to external persons without human review and verification review. In this respect, the generated summaries are treated pragmatically, not as official endpoints, but useful starting points that can be transformed into refined human-mediated documentation. Responsibility for producing accurate accounts is consequently distributed across both Plaito and meeting participants. As we discuss later, this distribution of agency and accountability extends beyond simple checking and editing of AI-generated summaries.

Through transcription and summarisation, the meeting is transformed from ephemeral events into persistent artefacts—working documents that give meetings material form.[1] This transformation has important implications for how collaborative work is organised. Documentation, for example, plays a significant role in how the organisation demonstrates professionalism to clients. At a practical level, automatically generated meeting records reduce the burden of notetaking during calls. But, by consequence of alleviating that burden, participants reported they were able to become more "*invested in the call*", listening more actively to the clients' concerns and requirements.

> "Historically, I would be on that call frantically taking notes and that can be a very distracted way to absorb information - or the cadence or pace of the conversation slows down because you're trying to catch up at all times where this actually allows me to be more of an active participant." P9 Senior Director

This was not simply about the pragmatic and cognitive concerns of achieving accurate information capture, but what it entailed for enacting a respectful relationship with the client. Not only were they better able to attend to the clients in the meeting without the distraction of notetaking, they were subsequently able to showcase their active engagement through improved documentation of the meeting outcomes.

> "I had a Rep that had a really high close rate and the reason that Rep was so good is he had the best notes of anybody on the team. Like you could tell when somebody spoke in his customer base, he truly listened to what they said. Well, now I have a computer that can do that. So, I can be even more active in a conversation. Like notes can be distracting." P10- Sales Rep

In addition to creating capacity for more active engagement with the client, we saw other important ways in which the LLM-enabled documentation and

[1] We have seen parallels to this in the context of other workplace communication where text-based conversations offered persistent material form to previously ephemeral spoken conversation (e.g. Erickson et al. 1999; Erickson et al. 2001).

summarisation of the meeting affected the in-situ organisation of conversation during the meeting. In this regard, the recording, transcription and summarisation of these meetings were not just neutral and mechanical representational outputs of the captured utterances. Rather, their presence had agency within the meeting, as participants adapted their utterances to the expectancies and operative procedures of the Plaito summarisation tool.

> "I find myself structuring the information more intentionally. Introducing the topic and providing the goal. Sometimes we'll jump into a meeting, forget to press record and then we'll press record and we'll go and introduce the meeting as if a new stakeholder just joined, 'we're talking about this and blah blah blah' and then we'll get it back into our discussion. I do find times where you would normally cross talk to a colleague, say like, 'oh, let's remember to do that', but you're just saying it for the sake of the recording. So, it's accessible to you later. I'll do that with JIRA ticket numbers. I'll find myself talking to the AI that's going to process the transcript." P5 Delivery Manager

The quote highlights how talk within these calls was not only reflexively accountable to colleagues and clients in the meeting, but also to the interpretive procedures that would ultimately render it legible. This "legibility" wasn't just the passive mechanical output, but the achieved work of making sense within this socially organised context of human-AI configuration. In changing their speech to accommodate the perceived operational contingencies of the recording and summarisation software, the participants simultaneously displayed and sustained the orderly properties of this new meeting setting—namely, that speaking was formatted to pass as machine-readable talk.

With experience, it was apparent that participants came to learn more about the system, its limitations and fallibilities and how to make it work well for them – adapting their behaviour in response to achieve the appropriate outputs. For example, in the following extract, the participant described how she would make small back-up reminders about what she considered especially important action commitments. These reminders were used as prompts during the meetings to verbally reiterate action commitments of both her and the client, specifically to render them salient to Plaito and ensure that they would appear in the subsequent summarisation.

> "I know that Plaito will send me the summary stuff, but sometimes they've missed things. So, I make sure I write it down on a separate screen, so I don't forget if a client has a question and I'll say yeah, I'll send you some documentation. Plaito sometimes doesn't catch that. And I don't want to say that I'm going to do something and not follow through on it, so I make sure

> I have a little list going on my other screen. Then, at the end of the call I'll say 'OK, to summarise, I'm going to do this, this and this. You are going to do this, this and this' – so that it [Plaito] reads the transcript. I make sure I say my action items multiple times throughout the call so I know for a fact it will be caught on the summary." P7 Business Analyst

There is an interesting sense in which the performative clarification and structuring of content for the benefit of the AI simultaneously served to signal its significance to others within the meeting. In making information legible to the AI system, participants were also rendering it salient to collaborators, subtly reinforcing its relevance to the ongoing work. In this way, the presence of recording and AI transcription technologies materially reshaped the dynamics of collaboration itself.

These arguments can be extended further to consider how assemblages of human-AI *intra-actions*[2] reconfigure the collaborative trajectories of organisational sense-making, decision-making, and practice. Importantly, these reconfigurations should not be understood as simply following a linear temporal sequence in which a meeting occurs first and summarisation subsequently acts upon it. Rather, the anticipated possibility of summarisation, and the opportunity for later agentic interrogation and restructuring of meeting accounts, transform the collaborative trajectory from the outset. An example articulated by P11, a talent acquisition specialist (TAS) illustrates this. P11 recounted a screening interview for a prospective job candidate that involved a trajectory of collaborative encounters between herself, the hiring manager, the job candidate and Plaito. Her responsibility was to produce a suitably structured account of the candidate to support review and decision making by the hiring manager. In this context, the online conversation with the candidate was shaped by the information requirements for her to make and evidence a recommendation to the hiring manager. Prior to the interview, the TAS and hiring manager had established, through Slack communications, a shared set of questions and information requirements that needed to be addressed during the screening process. Although P11 remained accountable to these agreed information requirements, she explained that the interview itself no longer needed to rigidly follow the predefined structure of the questions. Knowing that Plaito could later assist in restructuring and organising the discussion into a coherent account, she felt less constrained by the need to strictly adhere to the interview script in real time. As a result, she was able to engage with the candidate

[2] We draw on Barad's (2007) notion of *agential realism* which argues that phenomena are enacted through material discursive relationships or *intra-actions*. In this respect, the physical and social are not regarded as separate entities but are co-constituted in practice.

in a more fluid, responsive, and attentive manner, adapting the conversation dynamically as new topics and concerns emerged.

> "I think just really being in the moment with the conversation and…I kind of can see the questions and let them float into the conversation instead of having to: 'can you give me the answer for question one here's question two'. It's not so rigid I guess, so I'm able to work it into the conversation where they kind of don't know the questions being asked. So, I can actually listen to what they're saying and have a little bit more deep concentration with the conversation instead of having to type something out about an answer they said." - P11 Talent Acquisition Specialist

Following the interview, P11 then engaged with Plaito. The usual meeting summary output was not in a form that she needed to share with the hiring manager. As such she messaged Plaito in Slack:

> P11: "Based on the interview transcript [URL] can you answer the following questions [pastes list of questions previously agreed with the hiring manager]:"
> What is your location and availability for work
> Can you describe you experience with EPM and specifically with NetSuite EPM?
> How many years of experience do you have in financial planning, workforce planning and working capital implementation?
> What are your salary expectations?"
> Plaito: "Based on the transcript provided, here's what I can answer:

- [xxxx] is available to work in about one week. She is located in [xxxxx] and works 100% remotely
- [xxxx] has 5 years of experience working with EPM applications.
- …"

The bulleted list of responses provided a structured basis for her summary which she copied and edited before pasting into the candidate interview record. The final co-produced summary of the interview was of a suitable form for the needs of the hiring manager.

Additionally, she asked Plaito whether based on these answers, the candidate was considered suitable given the explicit requirements for the role. The generated recommendation from Plaito was not intended as any form of absolute decision outcome. Rather, it functioned as an evidence-based synthesis of the answers and their relationship to the role requirements. This helped P11

make sense of the information with a view to subsequent discussion with the hiring manager.

> "It's better with AI…because I know I can depend on it to find those answers. It really shortens my time, so I don't have to read through the transcription again or go back to look at the answers or find a specific word in the document. I know that AI can pull that a lot faster which saves me time. Before I used to do half hour blocks of interviews and have a half hour in between so I could finish my notes. [Now] I can do back-to-back interviews and just totally depend on AI to do all my summaries at the end of the day. It saves me a lot of time." P11 TAS

P11's account of being able to "depend" on the AI highlights the extent to which Plaito had become entangled within the material-discursive intra-actions shaping collaborative assessments of candidate competence and, ultimately, judgements of "fit" for the role. Interactions with Plaito were not simply extractive processes applied after the interview had concluded. Rather, they were grounded in, and shaped by, an earlier sequence of collaborative practices between the hiring manager and the talent acquisition specialist, through which hiring requirements and evaluation criteria had been negotiated, agreed upon, and documented. This list of questions influenced how the candidate interview was conducted, in the knowledge that it could be subsequently used as a prompt to Plaito to interrogate the interview transcript. The structured answers could then shape a response that could be looped back to the original information requirements, shared and rendered relevant to the hiring manager in their decision making.

Within this scenario, questions of responsibility and accountability are highly consequential, shaping whether candidates progress to subsequent stages of the hiring process. Importantly, these responsibilities are not simply delegated to Plaito. Rather, accountability is enacted as a collective accomplishment distributed across the participating actors. In part, Plaito's outputs become human-endorsed through processes of review, refinement, and incorporation into the talent acquisition specialist's communication with the hiring manager. However, accountability extends beyond the post-hoc verification and revision of AI-generated outputs. Equal attention must be given to the collaborative work that precedes and co-produces the material subsequently subjected to human–AI interrogation. The framing of hiring requirements, the structuring of interview encounters, and the generation of transcripts all shape the conditions under which AI summaries and analyses become possible and meaningful. In this respect, achieving accountable forms of AI-supported decision-making requires literacy not only in verifying AI outputs, but also in understanding and shaping the processes through which inputs to the AI are collaboratively produced.

### 4.2 Interrogating the Meeting: Collaborative Reconfigurations with AI-enabled "Ctrl-F" for Meetings

Accompanying the materialisation of meetings in document form was the opportunity to interrogate Plaito in search of specific information from the transcripts—what one participant described as a "*Ctrl-F for meetings*" (metaphorically alluding to the familiar short-cut used to search for specific keywords within a text-based document). In and of itself, this could be regarded as a valuable function of post meeting information retrieval, but its significance in the context of collaborative knowledge goes much further. Indeed, it is in the entanglement of these retrieval capabilities with the wider practices and apparatus of human engagement, recording, and transcription, that participation in collaborative knowledge work was further materially reconfigured. Let us consider an example recounted by one of the directors, discussing a core activity that they are required to do as part of winning new client business – namely responding to a Request for Proposal (RFP).[3]

> "Typically, people are hamstrung by the capacity of technical resource and so what [Plaito] has allowed for is to be more efficient with time. Let's say historically it takes six hours to review and respond to an RFP because of all the details and stuff that's there. And because of the different Subject Matter Experts (SMEs) we are looking at six hours by four folks across the team that's also pulling them away from other activities they could be a part of. A simple reflection of that is Bob. We've just got an RFP. Bob is my counterpart. 'What are your top ten questions of concern that you would need validated as a part of this?' So, I now have that Bob has spent five minutes giving me that list as opposed to two hours going through the document looking for that information. I've been empowered to have a tool to then capture that information as best as I can, whether that's on call with client and we're recording and transcribing so that Bob could then have that call later and go to that exact moment in time and hear his top concerns versus him actually being on that call. Or secondarily, just from capturing that from a 300-page document, you now have your control-F around the critical information so that you can get what you need. So, everyone has gotten time back in their day so they can be more impactful than some of the redundant tasks we do. We've elevated where we spend our time across the entire client acquisition to project delivery." P9 Senior Director

[3] A Request for Proposal is a formal document that organisations will use to invite proposals from service providers and vendors relating to specific project or service needs. The RFP will formally outline key information such as requirements and specifications, evaluation criteria, timelines and deadlines. These may be issued to multiple vendors at the same time who will then prepare their own proposal documents to be submitted as response to the RFP.

In this account, we see how, prior to the introduction of Plaito, responding to an RFP and preparing appropriate questions and responses for the client was an extremely time-intensive process. It typically required multiple subject matter experts (SMEs) to inspect large volumes of documentation and participate in lengthy meetings with both colleagues and clients, even when their expertise was only needed for relatively brief portions of the discussion. Underpinning this traditional process was a strong presumption of synchronous co-presence: SMEs needed to be available within meetings in order to support the rapid turnaround demands associated with RFP responses. This dependency on synchronous participation created significant coordination and articulation work. SME time allocated to one project inevitably reduced availability for others, meaning that scarce expertise had to be continually negotiated and coordinated across multiple concurrent project streams. However, the introduction of Plaito altered these collaborative dynamics by reducing the requirement for synchronous co-presence and enabling new forms of asynchronous and distributed collaboration. As a consequence, the organisation was able to substantially accelerate turnaround times for RFP responses while simultaneously reducing demands on SMEs and the coordination burden placed upon project teams.

Using responses previously provided by Bob, P9 was able to prompt Plaito to interrogate the extensive RFP documentation and extract relevant information in advance of meetings with the prospective client. Importantly, during subsequent client calls, P9 no longer required Bob's direct participation. Instead, P9 described verbally leaving what he referred to as "breadcrumbs" throughout the meeting conversation. These were interactional cues deliberately inserted into the discussion that would later become embedded within the meeting transcript. These cues enabled Bob to subsequently interrogate the transcript asynchronously through Plaito, allowing him to navigate directly to sections relevant to his own prioritised information needs without needing to attend the meeting itself.

This example illustrates how AI-supported interrogation of documents, recordings, and transcripts can reconfigure the very ontology of expert participation within collaborative knowledge work. By stabilising discussion into persistent artefacts, transcripts loosened the temporal coupling of collaborative activity and enabled work to be distributed across different times and locations. Yet these transformations involve more than simply enabling asynchronous work through AI-based transcription and retrieval. Coordination and articulation work are not eliminated; rather, they are reconfigured into new forms. Making this asynchronous collaboration possible required substantial preparatory work with SMEs to elicit the information that P9 could subsequently use to construct Plaito as a proxy for Bob during interrogation of the RFP materials. In turn, this groundwork enabled interactional signposting to occur during client calls, ensuring that the resulting transcript would remain intelligible and actionable for subsequent asynchronous review by absent experts.

The deliberate insertion of verbal "breadcrumbs" into meeting talk can therefore be understood as a form of transcript-oriented articulation work. Unlike more traditional forms of articulation work, transcript-oriented articulation practices involve interactional strategies through which participants deliberately shape spoken conversation so that its recorded afterlife can support future inspection, interpretation, and collaboration by absent colleagues. In contexts where transcripts are searchable, summarised, and algorithmically processed, such practices acquire infrastructural significance, as relatively small interactional cues can disproportionately shape how collective work is later reconstructed, delegated, and acted upon over time.

### 4.3 Meeting Content as Prompts in the Reconfiguration of Chat-Driven Authorship

Related reconfigurations of collaborative work were also seen in the authorship of documents. While the use of LLMs for facilitating content authoring is well established, this has tended to focus on their use by individual knowledge workers. Our observations, however, highlighted new ways in which content was "authored" through a collaborative network of material-discursive intra-actions involving multiple participants, tools and Plaito. Consider the following example, in which P11 has been asked by one of the senior directors to create some new LinkedIn messages to contact potential new talent for the organisation. To facilitate understanding, an overview of the scenario was as follows:

1. Initial call between P11 and Senior Director [Jack] to discuss this LinkedIn messaging requirements where Jack outlines content and style suggestions for the messaging.
2. Call is recorded and transcription produced by Plaito.
3. P11 engages in a conversational exchange with Plaito to extract the Senior Director's content requirements and use them to prompt Plaito to generate LinkedIn messaging content.
4. P11 shares Plaito outputs with Jack in Slack to solicit feedback.
5. P11 uses Jack's feedback to edit content in Google Docs.
6. P11 posts final drafts on LinkedIn.

Now let us look more closely at the exchange with Plaito in the context of this scenario:

> **P11:** "Summarise for what [Jack] is suggesting on the messaging we should use for LinkedIn [URL of meeting transcript]".
> **Plaito:** "Based on the transcript, [Jack] is suggesting the following approach for LinkedIn messaging:

1. Keep messages short and to the point. [Jack] emphasizes that people have limited attention, so messages should be brief (around 50 words or less).
2. Focus on "what's in it for me" from the recipient's perspective. The message should quickly convey the value proposition for the potential candidate.
3. For solution architects/consultants:
   - Mention the opportunity to transition into a full-time sales role
   - Include the potential On-Target Earnings (OTE) range
   - Ask if they're open to a brief conversation
4. For account executives/sales reps:
   - Highlight larger territories compared to their previous organizations
   - Mention the OTE range
   - Ask if they're interested in connecting for a brief conversation
5. Use language that creates curiosity or a fear of missing out. For example, for solution architects: "Given your past experience, I'm surprised you're not in sales."
6. Avoid including too much background information about the company in the initial message.
7. Make it easy for the recipient to respond with a simple yes or no.
8. The overall goal is to craft messages that are concise, relevant, and enticing enough to prompt a response or further conversation."

P11 then continues with a further prompt to Plaito:

**P11:** "With this information, provide me 3 examples of LinkedIn messaging we can use to reach out to potential candidates."
Here's the transcription of the text in your image:
**Plaito:** "Based on [Jack's] suggestions, here are 3 examples of LinkedIn messages you could use to reach out to potential candidates:
1. For a Solution Architect:
*Hi [Name],*
Given your experience as a Solution Architect, I'm surprised you're not in sales. We have strategic roles at Plative where you could leverage your technical expertise. OTE range: $150,000–$200,000.
Open to a quick chat about transitioning into sales?
2. For an Account Executive:
*Hi [Name],*
We're hiring for key Client Partner roles at [redacted]. Our reps consistently mention larger territories compared to previous companies. OTE potential: $200,000–$250,000.
Interested in a brief conversation about expanding your sales career?
3. For a Solution Consultant:

*Hi [Name],*
Looking to take ownership of the deals you support? We have opportunities at Plative for SCs ready to move into full sales roles. OTE range: $180,000–$230,000.
Are you open to a quick call to discuss?
Note: The OTE ranges used in these examples are fictional. Please adjust them to reflect accurate figures for your company. Also, remember to keep the messages short, focused on what's in it for the recipient, and easy to respond to with a simple yes or no."

On reviewing the suggestions from Plaito, P11 requests some refinements:

**P11:** "Expand the message a little more with what [Jack] is suggesting using the OTE as 160-260K. Keep it casual and active."

The exchange continued for a couple more prompted refinements. At this point, P11 copied the generated content into another Slack chat with the Senior Director with the accompanying message "*Thoughts?*" (on what has collectively generated through her engagement with Plaito*).* He responded in the Slack chat with some final suggestions:

**Jack:** "Short and sweet! I like it. Only change would be to have a clear call to action with no text after it. Make that the last thing they read.

Examples

- "How about a quick call to explore how you could turn your product knowledge into a high-producing sales role?"
- "What is the best way to get 15 min on your calendar to discuss this opportunity?"
- "Are you open to a brief conversation?"

No need to add "No pressure, just checking" etc."

P11 incorporated the suggestions before taking the content into LinkedIn for final edits.

We can see in this scenario is how the ontology of writing becomes fundamentally reconfigured through the coming together of human and AI agencies in the collaborative production of content. Particularly notable is how the initial verbal discussion of content and form between P11 and her Director (Jack) is first materially manifest, through a Slack-based prompt to Plaito, into a text-based summary of key requirements. Through this materialisation of the meeting, P11 effectively engineered a structured representation form that could

later be invoked within a follow-up prompt, forming the foundation on which subsequent AI-generated content was constructed. In this respect, the prompt itself embodied a form of collective agency distributed across P11, Jack, and Plaito. The subsequent prompt reflected how the initial output generated by Plaito did not adequately capture the intended voice and suggestions, leading P11 to request that the system expand further on what "[Jack] is suggesting." Once satisfied with the revised output, P11 reintroduced the Director's voice into the process by sharing the generated content within a new Slack conversation, explicitly acknowledging her use of Plaito and inviting further feedback. Jack's positive response effectively endorsed the draft, requiring only a minor addition before the content was considered suitable for posting.

The resulting content was therefore not simply authored in any conventional sense but collectively enacted through an evolving sequence of interactions between human participants and AI systems. Meetings, messages, comments, and prompts each contributed to shaping what the content ultimately became, foregrounding particular rhetorical forms and organisational values (e.g. short and succinct, enticing, limited attention, clarity of value) and marginalising others (e.g. not too much company background). Authorship here was not located in a single individual, tool or interaction. Rather, it emerged relationally through a distributed series of intra-actions across collaborating colleagues and Plaito.

### 4.4 "Remixing" Meeting Outputs into Project Documentation

Over and above documenting the meeting, the meeting transcripts opened new possibilities for the content to be "remixed" and transformed into new coordination artefacts, significantly reducing the burden of ongoing project documentation. An example of this was seen in the process of translating requirements from a client *Discovery Meeting* into trackable project management objects in Jira.[4] The key stages of the process are as follows:

7. **Discovery meeting**—all projects within the organisation are initiated by a discovery phase in which they engage the client in one or more discovery meetings to understand their specific requirements.
8. **Statement of Work**—based on these discovery meetings, a *statement of work* is produced outlining the high-level client requirements to be addressed by the project.
9. **User Stories Spreadsheet**—the high-level requirements are then formalised into a series of *User Stories* with associated "*Solution components*", "*Accept-*

[4] Jira is a software platform widely used in IT and software development projects (using Scrum or Agile methodologies) to track work, tasks, bugs, or other issues within a project. Within Jira, each "issue" or task is captured digitally as a "ticket" containing information such as task title, task description, person assigned to the task, status (e.g. to do, in progress, done), priority, additional comments/attachments for collaboration.

*ance Criteria*" and "*Level of Effort*" required – facilitating client engagement, resource planning and project execution. These components are documented in a project-specific *User Stories Spreadsheet* (see Fig. 3). The spreadsheet is a collaborative document authored by multiple Business Analysts and Project Leads across different project workstreams.

10. **Jira Tickets**—once the user stories spreadsheet is completed, it is submitted as an Agile request to a support team which translates the spreadsheet entries into associated Jira tickets. When the Jira tickets have been populated, the project lead will be notified at which point they can assign tasks to their team members to execute on the project.

Before the introduction of Plaito, generating user stories with the correct syntax and associated information was perceived to be time consuming and challenging. As P1 describes:

> "I used to be a BA before being a principal now and being a BA my nightmare was writing user stories. There were just so many of them and trying to get the syntax right, trying to think about all the possible acceptance criteria. But with Plaito we're able to cut documentation time in half." P1 Principal Consultant

P1 described how the process has now been simplified, by prompting Plaito (grounded in the "*Discovery*" meeting transcript) to generate the user stories and associated information to populate the *Project User Stories* spreadsheet. Each row of the spreadsheet corresponds to a single user story with columns for *Story Name, User Story, Acceptance Criteria, Solution Components* and *Level of Effort.*

> "So now I'm going to build all the User Stories and the Acceptance Criteria and then add the Solution Components and the Level of Effort for each item in this sheet. This is based off the 'Statement of Work' and based off the discoveries that we've had with the customer… So, when Plaito writes those user stories, we can update the spreadsheet based off the feedback from Plaito. So, I can pick an example here. [Prompt to Plaito] Write me a user story about setting up an email case in Salesforce. Include Acceptance Criteria and the Solution Components and…the Level of Effort for build in Salesforce. And it will try to get me some answers in Salesforce. Once it's done, it will say 'User story. As a customer service manager, I want to be able to set up email to case in Salesforce so that customers can easily and automatically convert cases allowing our support team to efficiently manage and respond to inquiries. This is the acceptance criteria, solution components, level of effort...' I would have to review it as a person to make sure

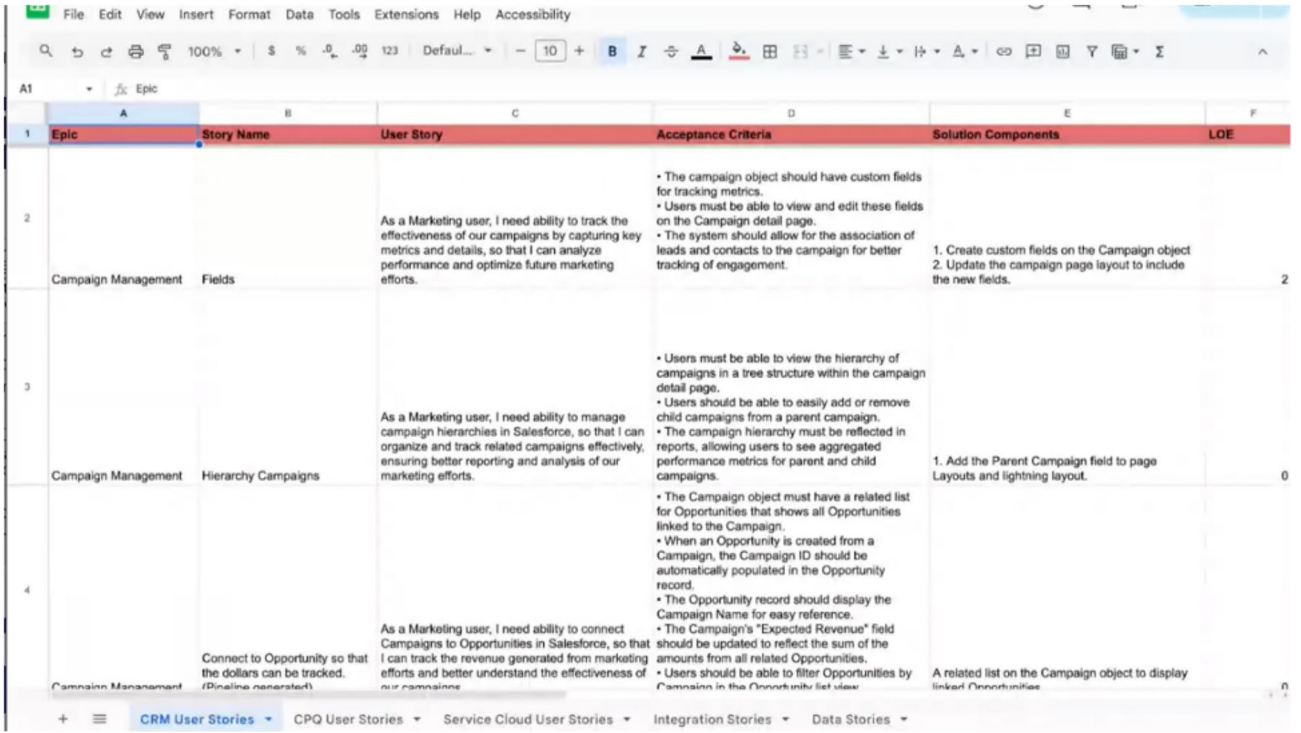


| | A | B | C | D | E | F |
|---|---|---|---|---|---|---|
| 1 | Epic | Story Name | User Story | Acceptance Criteria | Solution Components | LOE |
| 2 | Campaign Management | Fields | As a Marketing user, I need ability to track the effectiveness of our campaigns by capturing key metrics and details, so that I can analyze performance and optimize future marketing efforts. | • The campaign object should have custom fields for tracking metrics.<br>• Users must be able to view and edit these fields on the Campaign detail page.<br>• The system should allow for the association of leads and contacts to the campaign for better tracking of engagement. | 1. Create custom fields on the Campaign object<br>2. Update the campaign page layout to include the new fields. | 2 |
| 3 | Campaign Management | Hierarchy Campaigns | As a Marketing user, I need ability to manage campaign hierarchies in Salesforce, so that I can organize and track related campaigns effectively, ensuring better reporting and analysis of our marketing efforts. | • Users must be able to view the hierarchy of campaigns in a tree structure within the campaign detail page.<br>• Users should be able to easily add or remove child campaigns from a parent campaign.<br>• The campaign hierarchy must be reflected in reports, allowing users to see aggregated performance metrics for parent and child campaigns. | 1. Add the Parent Campaign field to page Layouts and lightning layout. | 0 |
| 4 | Campaign Management | Connect to Opportunity so that the dollars can be tracked. (Pipeline generated) | As a Marketing user, I need ability to connect Campaigns to Opportunities in Salesforce, so that I can track the revenue generated from marketing efforts and better understand the effectiveness of our campaigns | • The Campaign object must have a related list for Opportunities that shows all Opportunities linked to the Campaign.<br>• When an Opportunity is created from a Campaign, the Campaign ID should be automatically populated in the Opportunity record.<br>• The Opportunity record should display the Campaign Name for easy reference.<br>• The Campaign's "Expected Revenue" field should be updated to reflect the sum of the amounts from all related Opportunities.<br>• Users should be able to filter Opportunities by Campaign in the Opportunity list view | A related list on the Campaign object to display linked Opportunities | 0 |

**Fig. 3** User stories spreadsheet in preparation for conversion to Jira tickets

it aligns with what Salesforce can do because there's still a human part of it." P1 Principal Consultant

The output responses from Plaito are copied and pasted into the relevant cells in the spreadsheet. There is an acceptance that responses may not always be perfect and that there is a need for them to be reviewed and edited as necessary. Subject matter expertise of the Project Lead thus remains key to verifying that the Plaito-generated spreadsheet entries are correct—much in the same way as they had to do when user stories were populated by junior Business Analysts.

"So, if a BA did this and wrote the stories, I will still have to go in and manually adjust the level of effort if I think it should take lesser hours of time." P1 Principal Consultant

Once the entries are verified, the spreadsheet is then ready to be submitted as part of an Agile Support request workflow to convert them to Jira ticket entries.

"In Slack we have a couple of workflows. We have a support channel. So, we don't get to load this into Jira. We have a team that does that. Once the spreadsheet is ready to be loaded, you create a 'submit an agile support' request and it will run the workflow. It will ask you for the priority (medium, low). That way you can get it done faster. Put the description, attach the file, and somebody in the channel will let you know when the request is complete and then all those stories get loaded into Jira." P1 Principal Consultant

Central to these sequences is the extent to which collaborative knowledge workflows remain mediated through an established set of organisational coordination artefacts and processes. The spreadsheet, for example, is a pre-existing artefact used to coordinate the assembly, documentation, and handover of information across multiple teams and project leads, supporting collective review and validation before work progresses downstream. Such workflows reflect hard-won local knowledge and tacit practices developed over time, in which artefacts, including shared documents, ticketing systems, status boards, and informal channels such as Slack, act as scaffolding for mutual awareness, accountability, and the temporal organisation of work. The bottom-up "remixing" approach observed here, in which Plaito was used to transform information into forms appropriate for existing workflows, preserves and builds upon the significance of these artefacts. Rather than displacing established practices, this approach remains sensitive to the situated coordination routines already in place. As a result, users were able to appropriate the LLM in ways that made sense locally, adapting its outputs to the norms, constraints, and expectations of their specific teams, roles, and domains.

This also highlights how interactions with Plaito supported incremental, micro-level workflow improvements without requiring global consensus or formal process redesign that is often difficult to achieve and can be misaligned with everyday practice. By embedding LLM use within familiar coordination artefacts and workflows, adoption was facilitated rather than disrupted. Users could recognise immediate value while retaining agency over how the system was used, enabling a more adaptive and situated evolution of work practices. Over time, such micro-adjustments may accumulate into more substantial transformations; however, the notion of "remixing" underscores how these changes emerge from within practice itself, rather than being imposed through external mandates.

### 4.5 Documents as Meetings and Process Improvement

Thus far, we have shown how the materialisation of meetings through recording, automated transcription, and AI-based interrogation reshaped the production, sense-making, and accountability of collaborative work. As participants became more familiar with these capabilities and integrated them into their everyday practices, they began to appropriate the tools for the purpose of improving collaborative processes themselves. A particularly illustrative example was the introduction of new templated documents designed to support recurring organisational activities. These templates were not merely records of work, but were actively used to structure talk within meetings, with the explicit intention of enabling the subsequent population of these documents through AI-supported transcription and summarisation. In this context, P6 described how project teams regularly conduct *Design Review* meetings with clients. During these sessions, teams discuss candidate solutions developed to address client automation needs, articulating their respective advantages and limitations. To support this process, P6 explained how she prompted Plaito to generate a Design Review template containing appropriate sub-headings. These headings were specifically designed to capture and organise the outputs of such meetings, thereby providing a structured framework through which discussion could be guided in real time and later rendered into a coherent, standardised document.

> "What we've done is have several different interviews [meetings with solution designers & developers] which we record intentionally with the purpose of recording it and taking that information back to this template. So, this is the format that we would use for each of our Design Reviews. We give a little bit of information about the background, context, requirements [sub-headings], and then go into what different considerations we need. In this one, we had two different initial solutions that didn't work. So, what we're going to do later is review each solution and see why not. We essentially put together two different solutions and were able to speak from these

> interviews with our solution designers and developers to then be able to dictate into this document what we're going to do and why." P6 Consultant

The created template, which previously would have been a tool for post-meeting documentation, was now operating as a live, conversational resource during these meetings. The structure and subheadings of the document were used as a visual reference during these meetings to help organise topic sequencing, determine what should be discussed when, and framing whether the talk had achieved a level of adequacy for documentation. Again, participants in these meetings organised their talk to be "hearable" by the LLM agent as orderly, complete, and categorisable under the prospective headings of the Design Review document—"*Overview*, *Background*, *Context*, *Requirements*, *Solution Details*, *Considerations*, *Initial Solutions*, *Implementation Details* etc." The meeting talk was deliberately produced in the meeting with an orientation to its future recording and documentation. Participants made their contributions both to advance discussion during the meeting and talk the templated document/process into existence.

> "When I'm in the meeting it might be something that I don't understand but when I'm talking to my developers, I tell them, for the context of this recording and using a solution design template, tell me about [x]. We've adjusted our thinking based on how we're going to use the AI in the context of the recording. Now we'll be able to specifically use a template for the solution design and be able to lay out everything in that recording. And if I don't understand exactly what they're talking about, we can still speak to it in the context of the recording. Before using templates and recordings, writing this document, I would most likely try to understand through what our developers are talking about which is harder for me because I don't know anything about code. So, with using AI, they can tell me and the AI can pull the exact context of the language" P6 Consultant

The participants' methods here show that they are operating under a prospective accountability towards the AI-generated document. Their talk is produced not just to make sense in the moment, in the meeting, but so that it is suitable for future reading as coherent, template-fitted entries. Towards this end, they would do things such as reformulating each other's utterances for clarity, inviting specific articulation for the purposes of the recording and AI, and deliberately making otherwise tacit information explicit and visible to ensure they end up documented. Thus, the work of making talk documentable is itself visible, witnessable, and collaborative. In this regard, these practices exhibit the practical accomplishment of a *double reflexivity* in which the meeting participants first reflexively sustain the meeting's local coherence and structure, and second, reflexively sustain the meeting as an event for documentation. While technically

passive during the meeting, the recording and the templated LLM translation demonstrated agency over the meeting talk by exerting a structuring force that participants accountably oriented to.

Other participants discussed the role of Plaito in creating more standardised processes across the organisation that could be used for the generation of organisationally consistent artefacts and processes. A key example here was articulated by P14 (Head of Talent Acquisition). She described how the HR team had put together a *Career Framework* document. This document captured key job profiles within the organisation, associated responsibilities and expectations for development and progression at different levels of seniority. This framework was then used as the basis for grounding interactions with Plaito in ways that provided consistent HR-related processes and content across the entire organisation.

> "Our HR team has put together our career framework - a bunch of bullet points of what we expect in a senior consultant versus a consultant versus a business analyst. AI has made it incredibly easy to take our career framework and create job descriptions with it in a way that it's not just copy and paste. So that has taken a lot of time off people's plates because at the very least it gives us some bare bones to work with and we know for certain that it lines up with our career framework. I will take our career framework and say to Plaito: 'Based off of these bullet points or based off of these characteristics for senior sales force consultants, please create a job description for me'. It'll just pop something out – a beginning, responsibilities, qualifications, all of that stuff. I will share that with them [hiring managers]. We put it in a Google doc and then we collaborate together on it and ensure it's not completely out of whack." P14 Head of Talent Acquisition

This same framework document provided the basis for other aspects of the HR pipeline. P14 described how the framework was used to "ground" Plaito in the generation of subsequent interview questions that the hiring manager would be able to ask – in ways that were aligned with the career framework:

> "One way AI has helped in the later interview process is to create interview questions the hiring manager can ask based off the career framework and different core competencies we rate people on at Plative. That's been a game changer because lots of people don't know how to interview. It was a big thing for us before, trying to create more structured interview questions because theoretically we should be asking the same questions to every person to get an unbiased view. So, we went on this journey to create more structured questions that everyone is asking people. Before Plaito, I was sitting down with hiring managers, saying 'give me a list of questions' which was a disaster because they have to take up their time. Now I take the job description, and say 'based

> off these bullet points for senior sales force consultants, create interview questions that I might ask a candidate to assess if they are the right fit for this role and Plaito will just pop out questions" P14 Head of Talent Acquisition

By grounding Plaito in the career framework, the framework was no longer a static document to be consulted and interpreted in the performance of disparate HR activities (from talent acquisition through to subsequent performance assessment, learning and development). Rather, the framework was transformed into a more active coordination artefact that mediated across HR practices and stakeholders. Previously, the work of interpreting, translating and aligning HR-related intentions with the career framework involved repeated negotiation between stakeholders via Slack or online meetings to achieve actionable consensus. By using Plaito, the localised intent (create job advert, generate interview questions) could be translated into framework terms on demand and at the point of action. In this respect, much of the previous articulation work of interpretation and alignment across the different stakeholders could now be delegated to Plaito.

### 4.6 Augmenting Work of Situational Awareness Across Multiple Projects and Teams

Within the organisation, many roles, such as senior executives, project managers and capability leads, were responsible for multiple projects, resources and client engagements. In the course of their everyday work, achieving *situational awareness* across these projects is fundamental to successful performance. This is achieved through a complex weaving together of heterogenous practices to make overall sense of the project and determine appropriate action. These practices would include informal talk, chat-based retrievals of project status checks, review of project-related documents (meeting summaries, CRM notes, budget sheets etc.), memory invocation (e.g. recalling past conversations, decisions and client concerns), relevance calibration (e.g. importance, prioritisation and urgency) and anticipation of potential future issues before they arise. Situational awareness, was not just something attained but, rather, was an ongoingly and actively achieved practice. The methods for doing this were bound up in what we might conceive as the *managerial umwelt*[5] – the organisational world as it is perceived, structured and made actionable based on their specific needs and capabilities to access, sense and interpret organisational information. Prior to the introduction of Plaito, the *managerial umwelt* was bound up in a disparate set of artefacts, documents, databases, conversations, memories and various siloed systems. Such

[5] *Umwelt* is a term coined by Jakob von Uexküll (2010) The term refers to the subjective perceptual world of an individual organism; how that individual organism experiences and interprets the environment based on its sensory capabilities. The term is used more specifically in relation to technology and AI by James Bridle (2023) in *Ways of Being*.

fragmentation constrained what was visible, knowable and actionable for them. The previous methods and practices for producing situational awareness required them to navigate across these disparate sources and actively stitch things together. This was both time consuming and cognitively challenging. As P2 recounts:

> "As a manager I have probably 50 open projects on the go at any one moment. I have to stay on top of all of them. I also like to stay on top of all the projects going on in other parts of the business, that are related to me, but I don't want to be watching every single message. So, I put them on recap and I read the recaps probably daily. It's caused some problems because I've missed some stuff where people tagged me, but I have a good overview, so I have to oversee so much information. That recap and the summarisations of channels and stuff have given me a good way to kind of get through that all quicker… Slack is our front end to query all this information" - P2 Senior Director

With the introduction of Plaito to the Slack platform, we saw evidenced new ways to make such disparate organisational information more immediately visible, knowable and actionable. This was both through the project "short cut" buttons (which are essentially frequently required prompts—Project Status, Risks, Budget, Invoices, Sentiment, Meetings, Time Logs, Slack, Jira and Tracker)—and through more directed interrogation of Plaito. Through these interrogations, managers were able to engage in new forms of *aggregative surfacing* and *temporal stitching* – whereby Plaito would provide consolidated project-relevant views and summaries across different information systems and across different temporal periods of the projects. The manager's situated use of the agent's outputs changed how they were able to sense, sample and perceive their organisational world to achieve a new *managerial umwelt* where projects, clients, resources, deadlines, risks, and opportunities could be stitched together into continuously accessible, queryable and remixable forms.

> "If I want to ask regarding a specific project, I could click on 'status', 'risk', 'budget', 'invoices' - again it's using the data of that project in Salesforce and [there are] all these buttons - Status, Risk, Budget, Invoices Sentiment, Meetings, Time logs, Slack, Jira, Tracker et cetera - if I click on "Status" [button] it's going to reply asking Plaito to summarise the project status. After reviewing all project trackers budget, it comes through as this: 'Plaito has been grounded with the selected Project and related records for Accounts, Trackers, Releases, Risks, Invoices and CSATs… Based on the information provided, here's a summary of the [project name]". I'm using that to identify what's going on with different projects. If I'm going on a

> vacation for five days and I want a summary of what happened, I would say, 'Plaito, tell me what's happened in the last five days?'" P3 Program Manager

Further building on this idea of *managerial umwelt*, Plaito offered interesting new ways to sense and maintain situational awareness over different projects, clients and people. Of particular significance here was the ability to provide an indication and measure of project and meeting sentiment. Based on meeting transcripts, Plaito algorithmically generates a sentiment score for a meeting on a scale of 1 to 5, with 1 reflecting negative sentiment, and 5 reflecting very positive sentiment. This ranking appears in the meeting summary with a brief textual description to explain the basis of the score. For example, *"Sentiment: Positive (4 of 5) – The meeting had a neutral to positive tone with no strong expressions of frustration or disagreement. There was a focus on understanding processes and planning next steps, indicating a productive session."* (P3 meeting summary). As we can see here, Plaito's explanation offered only a high-level account of the sentiment. In and of itself, this may not fully capture the full nuanced understanding of the meeting dynamics. There were times also where Plaito would over-index on particular linguistic features of the conversation that could skew the overall score and summary in both positive and negative directions.

Such forms of automated algorithmic monitoring raise significant ethical questions regarding their practical deployment, particularly given the possibility of sentiment misrepresentation and contextual misreading. Recording meetings for the purposes of revisiting, interrogating, and remixing content is one matter; using those same recordings to algorithmically score and assess meeting sentiment is a far more delicate and contentious undertaking. Whether such assessments ultimately function as mechanisms of surveillance or, alternatively, as tools for oversight and collective reflection depends heavily on their organisational purpose, consequences, and situated practices of use.

Within the organisation studied here, sentiment analysis was not deployed as a mechanism for monitoring employee performance through forms of algorithmic surveillance. Rather, it functioned as a coordination device intended to support collective reflection and organisational sense-making across project teams and client relationships. Underpinning this use of sentiment analysis was the practical reality that senior management could not be "*in every place they need to be at any given time*" (P9) in order to oversee the full range of clients, projects, and evolving resource demands simultaneously. Consequently, managers were often unable to participate directly in all meetings occurring across their client portfolio, limiting their ability to maintain up-to-date situational awareness and respond proactively to emerging concerns.

What emerged in the study was how Plaito's sentiment analysis was used to help address these limitations. Meetings receiving a rating of 3 or below,

indicating negative sentiment, were automatically flagged in a leadership Slack channel[6] for the attention of senior management. Importantly, however, these alerts were not treated as objective or definitive representations of meeting dynamics. Staff remained aware that Plaito could sometimes over-index on particular words or phrases, producing misleading interpretations of conversational tone or intent. What mattered organisationally was therefore not the sentiment score itself, but the accountability process it initiated. Negative sentiment flags created demands for explanation and triggered escalation practices through which managers, staff, and, where appropriate, clients collectively reflected upon and interpreted the situation. In this respect, sentiment analysis functioned less as an automated judgement system and more as a mechanism for prompting further collaborative inquiry and situated sense-making.

> "These reports actually go to [my manager]. They will read it and will say P3, what happened in this meeting? The sentiment is 2. And then I will explain, yeah, because this happened, that happened, et cetera. So, it goes to the leadership." P3 Program Manager

A number of important behaviours arise from this. First, anticipation of escalation and enquiry by senior management lead the consultant to reflect and understand potential situational factors contributing to the score. These may be genuine factors inherent in the client relationship and situation or other signals that may have contributed to a flawed AI's flawed assessment.[7] Through this reflection, consultants could pre-emptively prepare response to potential queries from leadership. Second, seeing the scores in the Slack channel provided mutual awareness for all the leaders in near-real time. Third, it allowed the appropriate senior leader the opportunity to investigate.

> "I'm going to do my pre-work. I'm going to meet with the consultant, better understand their perspective of the situation. Clients similarly will have their perspective of the situation. I have to figure out where the truth actually lies, but also how you get towards a resolution. So, Plaito gives a referee's view of the issue and that's where I used objective versus what client feels they've experienced and what other individual feels they've experienced. So, it allows you to have a more empirical approach to a conversation before you jump into it. It also allows me personally to go into a call

[6] The sentiment scores and assessments are transparently available to all employees on a project a. Scores of 3 or below are openly known to trigger alerts. The alerts appear in a dedicated leadership slack channel to ensure it is seen and acted upon rather than to limit any transparency.

[7] There are parallels here with Eshan's (2024) notion of Seamful XAI in which any imperfections of the AI algorithms are embraced in the construction of meaningful explanations in the broader context of workplace collaboration.

> with context yet still seek to understand. It's given me a tool to actually navigate that conversation with a baseline. And I could do that immediately as opposed to waiting to get back to the client who's fuming about this issue…because I don't feel comfortable to have this conversation until I get on a call with my consultant who's on back-to-back calls and now four hours have gone by before I've addressed this burning issue. So, it just increases the chance of productive or successful interaction at a critical point." - P9 Senior Director

The above account highlights how such conversations can be difficult to navigate and therefore require considerable sensitivity in how they are approached. Importantly, algorithmic monitoring was treated not as a definitive judgement, but as a starting point for further inquiry and resolution. Sentiment analysis functioned primarily as a mechanism for raising awareness of potential issues, escalating concerns to senior colleagues, and providing what P9 described as "data to leverage", a provisional grounding upon which subsequent investigation, discussion, and intervention could be carefully constructed. Crucially, the purpose of these processes was not to evaluate employees, but to support collective sensemaking and coordinated responses to emerging client or project concerns.

A significant factor enabling these escalation processes to function constructively was the organisation-wide recognition of the potential fallibility of AI-generated signals. The possibility that problematic sentiment scores could result from limitations or misinterpretations within the AI system itself allowed conversations to be approached in a less confrontational manner. Rather than presuming employee fault or poor performance, participants could orient to the alert as one possible interpretation among several. In this respect, sentiment signals were treated as interpretable and contestable devices, whose value lay in prompting reflection and coordination rather than serving as authoritative assessments of employee performance. This highlights the sensitivities of these systems to the organisational culture and values in which they are deployed and used. In more authoritative organisational cultures, however, there may be potential risks of these signals being reified as an objective records of meeting sentiment, lending undue weight to inherently ambiguous interpretations. This may shift its role from an object of enquiry to an evaluative metric, potentially constraining how people express themselves and how interactions are judged. Indeed, we see indications of employee evaluation awareness even within the more progressive culture of the current organisation. Even though the culture did not explicitly frame the system as a tool for employee monitoring, participants reported consciously adapting their behaviour during meetings in ways intended to produce more positive sentiment scores. Notably, this behavioural orientation persisted despite widespread awareness among staff of the potential inaccuracies and fallibilities of the AI system itself.

> "Yeah, I it's like hard because I don't think it's like super accurate. But I like to just see what my score is. I always try to get a five out of five because I just think it looks good… Yeah, senior leadership sees it. So, I always just try to get a high score." - P7 Business Analyst

### 4.7 Meeting Preparation and Context Switching Across Multiple Client Relationships

The nature of consultancy work means that many in the organisation are involved in numerous projects whereby work comprises intermittent engagements with multiple clients that may play out longitudinally over time. For example, one of the Business Analysts (P7) described how she was working collaboratively on approximately 15–20 projects across different clients. For each project, there were requirements to complete small discrete pieces of work intermittently distributed out over weeks and months. Given these commitments, the day-to-day organisation of her work entailed context switching across several different collaborative engagements bringing enormous challenges and required considerable work to manage:

> "I am a part of about 15 to 20 different projects right now, so I have to do a lot of context switching between my clients to think about 'Last time we met, we talked about this, these are my takeaways.' It's a lot to keep track of. So, I can talk to Plaito about my project and say 'Hey, when was the last time I talked to this client about this subject' and it'll say 'here are the meetings in which this topic was ever brought up' and I can ask questions 'OK did they ever decide what to do about this specific decision' and it would be 'Yes, it was decided in this call. Here's the chat transcript'. Or 'no, it was determined that we talked about this later'. Because you don't want to go to a client and ask them the same things over and over. They get annoyed. So having Plaito to ask those questions to is great, without having to rewatch recordings for hours, trying to make sure I'm finding the right recording, and being able to get the information I need quickly is helpful" - P7 Business Analyst

What P7 refers to here is the importance, within consulting work, of attending to the sequential organisation of client interactions over time. Across these evolving trajectories of engagement, P7's ability to recall and demonstrate knowledge of prior encounters, discussions, and unresolved questions functioned as an important means of maintaining the recognisable orderliness of professional relationships. From one interaction to the next, remembering, and crucially, displaying that remembering, was not merely a cognitive accomplishment, but a form of relationship work. Such relationship work involved accountable practices

through which participants oriented to the ongoing relevancies of the client relationship, framing interactions as part of a continuous collaborative history rather than as isolated and disconnected meetings. P7 described how this remembering was actively displayed within current conversations through explicit indexical references to matters discussed in earlier encounters. As she explained, failing to acknowledge or recall previous discussions risked being interpreted as inattentive, unprofessional, or even disrespectful. Conversely, the ability to weave prior conversations into present interactions enabled consultants to perform continuity, care, attentiveness, and competence in ways that were central to demonstrating professionalism to clients.

P7 further highlighted how this relationship work of recalling and reintegrating previous discussions into ongoing client conversations had historically been both difficult and time-consuming. Prior to the introduction of Plaito, preparing for client meetings often involved extensive effort locating relevant meetings, searching through recordings, and reviewing fragmented notes in order to reconstruct the history of prior discussions and decisions. With the introduction of Plaito, however, this preparatory work became substantially reconfigured. Rather than manually reconstructing interaction histories, P7 described being able to directly interrogate Plaito regarding previous client discussions, asking whether specific topics, decisions, or concerns had been raised previously, and when those discussions had occurred. For example:

> **P7**: "Can you show me the summary notes from the meeting from 2024-07−26?"
> **Plaito**: "Sure, here is the summary information provided for the MCAE[8] Project Planning Meeting that took place on 2024-07−26…
> **P7**: "Can you summarise any meeting where MCAE or Pardot[9] was brought up?"
> **Plaito**" Here are some of the key meetings where MCAE or Pardot was brought up"
> [lists meeting titles with links to transcripts].
> **P7** "What do they want to use MCAE or Pardot for"?
> **Plaito**: "Based on the context provided in the meeting summaries it seems that AcmeCo were looking to implement either Pardot or MCAE to address some key needs…"

The exchange with Plaito highlights some of the challenges in establishing the necessary context for an upcoming meeting. The interaction begins with a loosely

[8] Marketing Cloud Account Engagement (MCAE) is a B2B marketing automation platform by Salesforce.
[9] MCAE was formerly known as Pardot.

formulated information goal on the part of P7. The request for summary of the MCAE project planning meeting offers initial context but doesn't fully satisfy the emergent information need about MCAE and Pardot. Additional prompting by P7 then enabled a more targeted exploration of other meeting summaries with a focus on MCAE and Pardot. From this, P7 was able to further formulate a more precise query about the Client's intended use of MCAE and Pardot. Through this interactive interrogation of Plaito grounded in the previous client meeting transcripts, vague recollections guiding context retrieval can quickly be transformed into more precisely targeted queries to obtain the key pieces of context needed for the subsequent client meeting. In turn, through quickly obtaining this contextual information, they were better equipped to display professional competency and perform professional relationship work through the appropriate indexical reference to previous discussions.

Related behaviours were apparent too in the work of establishing new client relationships and leads. For the Business Development Managers, a significant proportion of their work involves speaking with multiple new clients. Each day, there will be calendar entries for meetings with new clients who have responded to initial outreach efforts by the BDM. Some of these calendar entries may have been initiated by the prospective client who wants to know more information about the offering. This can come as a surprise for the BDM, leaving them with very little context for the call.

> "Usually very cold calls, you know, someone who I've sent an e-mail to and they've just booked a meeting. So, they've seen my e-mail and they book a meeting with me via my automatic calendar link. And then it's like, I have meeting, who booked me - sometimes it can be a bit of a surprise." P8 BDM

Again, as a consultancy organisation, and in particular when prospecting for new business, presenting a professional persona is a critical part of that relationship building process that in turn can contribute to winning new business. This entailed call preparation work to understand the client, the company, their role in the business, and what they care about. Such understanding allowed them to approach the conversation in the right way and display professionalism. Prior to the introduction of Plaito, such preparation work was time consuming and needed explicit structure to the daily schedule to accommodate it. But as P8 described, this would often leave them feeling rushed.

> "I feel a little less rushed to be prepared for meetings. I usually would give myself a 15–20-min gap between any customer meeting so that I could go and find this person, find the company. Now 5 min before the call, I can hit the "Call prep" button and do a quick review of everything about this per-

> son and head into the call with confidence to speak to some key pieces… Obviously, it's not going to be everything, but it's enough to have a warm conversation. I can just focus on talking to the clients because I can get that information so much faster. I can speak to more people, in a more meaningful way, in a shorter period of time." P8 BDM

With the introduction of Plaito, they could manage this context switching and call preparation much more efficiently. The dedicated "call prep" button is associated with each calendar entry, situating its capability where it is needed in the context of their collaborative task flow – it is at the right place at the right time, grounded in the context of the work. Of significance here is how such efficient aggregation of the relevant information about upcoming call participants provides enough context to enable them to engage in "warm conversations" and talk in a more "meaningful way" with these clients – all important elements of professional business development encounter. These efficiencies make the transition between calls much less stressful, reducing tensions between the everyday realities of post-call activities of one call and the pre-call activities of another:

> "So many times it just takes out the work that I need to do, especially because I'm running for the next call while trying to send a summary for the previous call. It's just faster to get things done than have to spend that time trying to figure out the different pieces." P1 Principal Consultant

## 5 Discussion and Conclusion

We have presented a study of the evolving practices of knowledge workers within an IT services organisation following the introduction of an organisationally tailored LLM agent embedded within the enterprise systems, tools, and documents constituting the organisation's knowledge infrastructure. As enterprise deployment of such systems accelerates, there is an increasing need to understand how they transform the structure and organisation of collaborative knowledge work. Addressing this challenge required moving beyond dominant accounts of LLM use centred on individual productivity gains, bounded task performance, or narrowly defined human-AI teaming dynamics. Instead, our research adopted a situated practices perspective on emerging forms of LLM-mediated collaboration through the in-situ study of an enterprise-specific agent embedded within real organisational workflows over time. By attending to these situated practices, we reveal how value emerges not simply at the level of individual task completion, but through collective sense-making, the redistribution of effort, and the reorganisation of collaborative work itself.

Rather than merely supporting discrete activities such as drafting or summarisation, our findings demonstrate how infrastructurally embedded LLM agents more fundamentally reconfigure the relationships between meetings, documentation, and asynchronous collaboration. These reconfigurations produced important shifts in organisational practices of shared sense-making, documentation, accountability, awareness, articulation, and coordination. Meetings became transformed into persistent and queryable artefacts through which collective knowledge and interaction histories could be recovered across both people and time. At the same time, collaborative outputs increasingly emerged through extended chains of human-AI interaction rather than through isolated acts of authorship. Importantly, workers adapted their own communicative and interactional practices in response to the presence of the system itself, deliberately structuring talk, foregrounding context, and shaping interactions in ways that rendered subsequent transcription, summarisation, and retrieval more meaningful and actionable. In this respect, the LLM agent did not merely support existing work practices but actively participated in shaping and reorganising them.

Taken together, these findings advance understandings of LLMs as organisational infrastructure rather than standalone productivity tools. They offer a richer account of how generative AI becomes embedded within, and increasingly constitutive of, collaborative knowledge work, while also foregrounding the new forms of accountability, coordination, and organisational risk that accompany such infrastructural embedding. More broadly, this perspective provides an important foundation for the future design, evaluation, and governance of AI systems operating within complex organisational settings.

Apparent throughout the fieldwork was that the utility and influence of these systems could not be understood in isolation from their embedding within the wider organisational ecosystem of enterprise tools, infrastructures, and knowledge practices. A particularly important example was the explicit integration of the agent within the Slack communication and collaboration environment. This integration enabled the agent to leverage pre-existing organisational structures for client, project, and team management that already played a central role in shaping collaborative work practices. Importantly, these structures within Slack extended into a broader ecosystem of enterprise, customer, employee, and project management systems, creating additional pathways through which information, context, and organisational relevance could be mobilised. It was this infrastructural embedding that enabled the agent's responses to become grounded in project-specific and client-specific ways rather than operating as generic or context-free outputs. We conceptualise this through the notion of *agential ubiety*, foregrounding the importance of the agent being situated in a particular place and time. The concept of *ubiety* features in Buxton's (2019) discussion of ubiquitous computing, smart devices, and device ecologies, where he argues that intelligence and utility do not reside inherently within devices themselves. Rather,

smartness emerges through the fit, placement, and appropriateness of technologies within particular contexts and through their interrelationship with wider ecosystems of tools and artefacts. Extending this argument, the concept of agential ubiety highlights how the value and agency of LLM systems are amplified through their careful and contextually appropriate embedding within existing organisational infrastructures. Their effectiveness arises not simply from model capability alone, but from their situated integration within the communicative, documentary, and coordination practices through which organisational work is already accomplished.

Of further significance within the study was the material transformation of meetings through processes of recording, transcription, and summarisation. This echoes early discussions by Erickson and colleagues concerning text-based computer-mediated communication systems that functioned simultaneously as sites of interaction and as persistent forms of documentation (e.g. Erickson et al. 1999; Erickson et al. 2001; Erickson and Herring 2004). In our analysis, we similarly observed how online video and hybrid meetings became materially stabilised as persistent artefacts. However, understanding the implications of these systems for collaborative knowledge work requires moving beyond the notion of passive documentation of pre-existing practices. Rather, our findings reveal the more active agency of these systems in shaping collaboration itself, as workers increasingly adapted their behaviour in anticipation of the prospective capabilities afforded by the persistent meeting. In the first instance, participants modified aspects of their communicative practice in ways that rendered their contributions more legible to the AI system, while also identifying opportunities where feedback mechanisms could further support such legibility. These adaptations produced tangible benefits, particularly by enabling participants to devote greater attention to clients during meetings, thereby supporting forms of trust-building that are especially important within service-oriented work.

Importantly, this transformation involved more than simply offloading the cognitive burden of note-taking onto an automated system. As workers became increasingly familiar with the future possibilities for interrogating, remixing, and repurposing persistent meeting artefacts, they began adapting in-the-moment interactional practices to better align with the social and relational demands of encounters while still satisfying the standardised documentation requirements of organisational work. For example, interview conversations no longer needed to rigidly adhere to predefined question structures because participants anticipated that structure could subsequently be reconstructed through agent-based interrogation and remixing of meeting transcripts.

Fully understanding the impact of these tools on knowledge work, and how we might effectively design for them, requires greater attention to their role across collaborative project trajectories. These systems can be understood as a dynamic and personalised form of "meeting bridging" (Wang et al. 2024) linking

the synchronous activity of meetings with broader forms of synchronous and asynchronous collaboration, documentation, and coordination. Importantly, the effects of these collaborative trajectories must be considered bi-directionally, encompassing both retrospective and prospective influences on the situated organisation of work. As meetings became materialised through recording and transcription, the ability to search and interrogate them through the LLM agent enabled transcripts to be remixed into new forms of collaborative and coordination artefacts. Frequently used organisational templates, for example, were actively brought into meetings to structure discussion in ways that enabled the agent to more effectively "fill in the blanks" later. Participants effectively talked documents into existence, asking questions designed to "*lead the [collaborative] witness*" such that subsequent transcription and prompting could reconstruct organisationally meaningful outputs. These practices were grounded in existing organisational templates and documentation standards, again demonstrating how the utility of LLM agents is closely tied to their capacity to leverage and amplify familiar organisational coordination artefacts. At the same time, participants also developed entirely new templates and workflows to exploit the emerging possibilities afforded by persistent, queryable meetings. Through the appropriation of familiar enterprise productivity tools and prompting practices, workers created improvised semi-automated workflows tailored to the situated demands of their own work contexts. Particularly significant here was the bottom-up character of these practices. Rather than workflows being imposed solely from above (although formal workflows still remained important) workers engaged in highly localised and improvisational forms of infrastructural adaptation using familiar tools and organisational resources.

Other important enablers of collaboration across project trajectories related to practices of re-establishing project context and maintaining situational awareness. These practices were facilitated through what might be described as a new form of *managerial umwelt*, enabled by the LLM agent's ability to aggregate and temporally surface information drawn from disparate organisational systems, including customer, employee, and project management platforms. This capability supported new ways of "knowing the organisation" that would otherwise have required extensive cross-team articulation work. Particularly salient were shortcut prompts encoding commonly required aggregations of project status information. Access to these consolidated views quickly became embedded in everyday practice, especially for managers coordinating activity across multiple teams and projects. More broadly, the findings suggest the importance of enabling organisations to construct such aggregations around the repetitive and situated features of their own work practices. Beyond these standardised aggregations, situational awareness was further enhanced through more contextually tailored forms of *interrogative aggregation*, enabling workers to surface relevant information across otherwise siloed organisational systems in response to immediate

collaborative needs such as pre-meeting preparation, project coordination, or client-facing relationship work.

At the same time, these emerging forms of expanded organisational awareness also introduce important epistemic risks. LLM agents may privilege forms of information that are most legible to computational documentation, retrieval, and summarisation processes. In doing so, they risk conferring unwarranted epistemic authority upon particular representations while flattening uncertainty, ambiguity, or alternative perspectives in ways that may reproduce existing organisational power structures. Within the organisation studied here, these risks were partially mitigated through a culture of interpretive reflexivity in which LLM-mediated outputs were treated as provisional resources for sense-making rather than definitive accounts of record. This enabled the systems to support managerial and collaborative work without erasing the negotiated and contestable character of organisational knowledge. Participants also demonstrated an evolving literacy regarding how to make themselves legible to the agent and how to critically interpret its outputs. However, these practices and benefits may not transfer straightforwardly across other organisational settings. In more formal hierarchical environments, or contexts characterised by stronger accountability pressures and tighter temporal constraints, similar systems could encourage epistemic shortcuts in which AI-generated accounts become prematurely stabilised as authoritative representations. These may increase the stakes of accuracy and reinforcing auditability. In such settings, participants may orient to AI-generated records not only as aids to coordination, but as potential instruments of evaluation, leading to more cautious or strategic forms of communication. Similarly, while the distributed organisation examined here exhibited a pragmatic, experience-based trust in AI outputs, hierarchical contexts may impose more formal governance structures that constrain how and where such systems can be used. As a result, the introduction of LLM agents may amplify existing organisational tendencies, enabling fluid, adaptive collaboration in flatter, digitally mature environments, while reinforcing standardisation, visibility, and control in more bureaucratic settings. Such dynamics raise important questions regarding organisational decision-making, accountability, and the governance of AI-mediated knowledge work, warranting continued empirical attention.

Of further significance were the more innovative analytical capabilities emerging around materialised meetings, particularly sentiment analysis. These developments point toward new ways in which collaborative activity can be analysed, represented, and incorporated into the negotiated order of organisational work. Earlier CSCW research has long argued that relatively simple abstracted representations of activity can support socially translucent forms of collaboration when interpreted contextually (e.g. Erickson et al. 1999; Erickson et al. 2001; Kaye et al 2005; Kaye 2006; Nardi et al. 2000; O'Hara et al.

2014). The capacity of LLM agents to extract and analyse communicative characteristics from workplace interactions extends these possibilities, opening new avenues for understanding and coordinating collaborative activity through lightweight socially translucent representations.

Within this organisation, algorithmically generated sentiment representations reconfigured articulation work around the surfacing, interpretation, and management of collaborative breakdowns. Yet these practices remain deeply contingent upon organisational culture. As a medium-sized and culturally progressive SME characterised by consensual and transparent recording practices, sentiment signals were framed as lightweight coordination cues inviting collaborative interpretation of interactional difficulties, emotional strain, or emerging client risks. Treated as provisional and situated signals, and coupled with human oversight, these representations supported a negotiated order in which meaning remained collectively constructed rather than algorithmically imposed. In this context, sentiment-driven escalation practices were not viewed as replacements for managerial judgement, but rather as practical mechanisms for scaling noticing, sense-making, and coordinated response across distributed sociotechnical infrastructures.

Nevertheless, important open questions remain regarding how such systems might operate across other organisational cultures, sectors, and power relations. In less progressive environments, or within contexts characterised by adversarial and compliance-driven management regimes, similar technologies could readily be appropriated as instruments of surveillance and algorithmic management. Sentiment scores may become reified as objective indicators of employee performance or mobilised without adequate sensitivity to interactional nuance, client behaviour, or emotional labour. The same abstractions that create organisational value can simultaneously flatten social complexity, amplify power asymmetries, and undermine trust. As earlier CSCW debates around awareness and monitoring have argued (e.g. Gaver et al. 1992), these systems are best understood not as neutral analytics but as socio-technical interventions shaped by organisational norms, policies and existing monitoring regimes. Given growing concerns surrounding algorithmic management (e.g. Gray and Suri 2019; Moore 2018), there is a pressing need for informed and balanced approaches to the deployment of such systems as coordination technologies. Effective implementation requires foregrounding transparency, contestability, collective interpretation, worker voice, and careful consideration of organisational readiness and cultural fit. Such conditions are essential for preserving the negotiated and accountable character of collaborative work.

Finally, the findings point toward more fundamental ways in which LLM-mediated interaction can reconfigure workplace collaboration itself, reducing coordination burdens while enabling more efficient use of expert resources. In several cases, information elicited from subject matter experts was used to

structure client meetings conducted in their absence, with verbal "breadcrumbs" deliberately embedded within interaction so that absent experts could later interrogate transcripts asynchronously through the LLM system. These practices facilitated a shift away from the burdens of synchronous expert participation toward more distributed and asynchronous forms of engagement, reducing both time commitments and coordination overheads associated with collocated meetings. More broadly, such practices suggest important opportunities for rethinking how synchronous and asynchronous, remote and collocated forms of work may be organised within knowledge-intensive organisations.

In conclusion, this study has shown how the collaborative significance of organisational LLM agents emerges not simply from their standalone technical capabilities, but from their infrastructural embedding within organisational ecosystems of tools, documents, and practices, and from the situated ways in which workers appropriate them over time. By studying an enterprise-specific agent in situ, we have demonstrated how these systems become entangled within trajectories of collaborative work, reshaping meetings, documentation, coordination, accountability, and organisational awareness in ways that are collectively accomplished and culturally contingent. Adopting a situated practices perspective has enabled us to foreground not only the benefits of these systems, but also the new forms of risk, dependency, and organisational reconfiguration they introduce. For CSCW, these findings reinforce the importance of understanding organisational LLM agents not simply as tools or teammates, but as forms of sociotechnical infrastructure.

The findings point toward several important considerations for the organisational embedding of these systems, foregrounding infrastructural fit, interpretive support, and collaborative practice rather than individual interaction alone. First, organisational LLM agents should be designed for deep and legible integration within existing enterprise ecosystems, enabling them to leverage familiar artefacts and organisational structures such as project hierarchies, document templates, communication channels, and coordination practices. Such integration supports the emergence of value through continuity with established collaborative workflows rather than through the imposition of parallel AI-centred processes. Second, these systems should be designed and governed in ways that preserve the provisional and negotiated character of LLM-generated outputs, foregrounding transparency, contestability, and collective interpretation. Doing so can help mitigate risks associated with the reification of AI-generated summaries, analyses, and representations as authoritative organisational accounts. This points further to the value of LLM-based tools not residing in their technical capabilities alone, but in how they are taken up, interpreted, and embedded within existing organisational cultures and practices. Organisational norms around trust, accountability, and communication fundamentally shape whether these systems function as enablers of fluid collaboration or as instruments of control and formalisation.

Third, organisational deployments should support bottom-up appropriation practices, enabling workers to create, adapt, and share prompts, templates, and semi-automated workflows that align with the situated demands of their own work contexts rather than imposing rigid top-down procedures. Fourth, both the design and analysis of these systems must attend more carefully to the temporal and trajectory-based character of collaboration, including the bi-directional relationships between synchronous meetings, asynchronous coordination, documentation, and project work. Developing a richer understanding of these temporal dynamics opens important opportunities for reconfiguring organisational workflows around more flexible relationships between synchronous and asynchronous forms of collaboration.

As organisations continue to grapple with these transformations in their everyday practices and structures, there remains a strong need for sustained empirically grounded research capable of informing the development of more responsible and effective AI-mediated work environments. The findings highlight several important directions for future CSCW research. First, there is a need for longitudinal and comparative studies examining how infrastructurally embedded LLM agents evolve across different organisational cultures, sectors, and power relations, particularly to better understand how practices of documentation, accountability, coordination, and organisational sense-making develop over time. Second, further research is needed to investigate how different approaches to governance, transparency, and worker agency shape the interpretation and appropriation of LLM-mediated artefacts, especially in contexts characterised by greater organisational hierarchy, worker precarity, or regulatory pressure. Third, significant opportunities remain for studying the micro-interactional consequences of LLM presence within collaborative settings, including how workers adapt talk, timing, gesture, repair practices, and interactional organisation in anticipation of AI recording, analysis, and future retrieval. Fourth, there is considerable scope for further work examining how algorithmic representations, including summaries, aggregations, and sentiment analyses, become socially translucent, contested, resisted, or institutionalised within everyday organisational practice, and how such representations reshape articulation work, managerial awareness, and organisational accountability. Finally, as organisations increasingly deploy generative AI agents at scale, CSCW research must continue developing theoretical and methodological approaches capable of understanding LLMs not as discrete tools, but as evolving infrastructural actors whose effects unfold across trajectories of work, systems, relationships, and time.

## Author Contribution

Research design was conducted collectively by MM, KOH and MK. Primary research was conducted by KOH and MM supported by technical and

organisational guidance of MK. Initial qualitative analysis was conducted by MM and KOH followed by additional collaborative interpretation and refinement by MK. The final manuscript was written primarily by KOH and MM with collaborative review and refinement by MK.

## Funding

No funding was received in support of this work.

## Data Availability

The authors do not have permission to make the interviews and other data from the study available.

## Declarations

## Competing interests

The authors declare no competing interests.

## Ethical Approval

Then study was approved by the University of Bristol Research Ethics Board (ref 2024–22491-24550).

## Consent to Participate

All participants orally consented to take part in the study.

## References


Asthana, Sumit, Sagi Hilleli, Pencheng He, and Aaron Halfaker. 2025. Summaries, highlights, and action items: Design, implementation and evaluation of an LLM-powered meeting recap system. *Proceedings of the ACM on Human-Computer Interaction* 9 (2): 1–29. https://doi.org/10.1145/3711074.

Barad, Karen. 2007. *Meeting the universe halfway: Quantum physics and the entanglement of matter and meaning*. London: Duke University Press.

Bødker, Susanne. 1989. A human activity approach to user interfaces. *Human-Computer Interaction* 4 (3): 171–195. https://doi.org/10.1207/s15327051hci0403_1.

Brachman, Michelle, Amina El-Ashry, Casey Dugan, and Werner Geyer. 2024. How knowledge workers use and want to use LLMs in an enterprise context. *Extended abstracts of the ACM CHI Conference on Human factors in computing systems (CHI EA '24)*, 189: 1–8. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3613905.3650841.

Braun, Virginia, and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3 (2): 77–101. https://doi.org/10.1191/1478088706qp063oa.

Bridle, James. 2023. *Ways of being animals, plants, machines: The search for a planetary intelligence*. New York: Penguin Books.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2025. Generative AI at work. *Quarterly Journal of Economics* 140 (2): 889–942. https://doi.org/10.1093/qje/qjae044.

Butler, Jenna, Sonia Jaffe, Nancy Baym, Mary Czerwinski, Shamsi Iqbal, Kate Nowak, Sean Rintel, Abigail Sellen, Najeeb Abdulhamid, Judith Amores, Reid Andersen, Kagonya Awori, Maxamed Axmed, danah boyd, James Brand, Georg Buscher, Dean Carignan, Martin Chan, Adam Coleman, Scott Counts, Daepp Madeleine, Adam, Fourney, Daniel Goldstei, Andy Gordon, Aaron Halfaker, Javier Hernandez, Jake Hofman, Jenny Lay-Flurrie, Vera Liao, Siân Lindley, Sathish Manivannan, Charlton Mcilwain, Subigya Nepal, Jennifer Neville, Stephanie Nyairo, Jacki O'Neill, Victor Poznanski, Gonzalo Ramos, Nagu Rangan, Lacey Rosedale, David Rothschild, Tara Safavi, Advait Sarkar, Ava Scott, Chirag Shah, Neha Shah, Parikh, Teny Shapiro, Ryland Shaw, Auste Simkute, Jina Suh, Siddharth Suri, Ioana Tanase, Lev Tankelevitch, Adam Troy, Mengting Wan, Ryen White, Longqi Yang, Brent Hecht, and Jaime Teevan. 2023. Microsoft new future of work report 2023. *Microsoft Research Tech Report*, MSR-TR-2023–34. https://aka.ms/nfw2023

Butler, Jenna, Michaela Vorvoreanu, Rebecca Janssen, Abigail Sellen, Nicole Immorlica, Adam Troy, Advait, Sarkar, Alex Farach, Alex Chouldechova, Alexandra Olteanu, Alexia Cambon, Arjun Radhakrishna, Asta Roseway, Ben Zorn, Brent Hecht, Daniel Goldstein, Dhruv Joshi, Ed Cutrell, Emre Kiciman, Gonzalo Ramos, Gustavo Soares, Hanna Wallach, Ian Drosos, Jack Williams, Jacki O'Neill, Jake Hofman, Jamie Teevan, Javier Hernandez, Jennifer Wortman Vaughan, Jina Suh, John Tang, Justin Edwards, Kalika Bali, Kori Inkpen, Krishna Madhavan, Laylah Bulman, Leon Reicherts, Lev Tankelevitch, Longqi Yang, Martez Mott, Millicent Ochieng, Mercy Muchai, Nancy Baym, Najeeb Abdulhamid, Nicolai Marquardt, Ken Hinckley, Michael Bentley, Dave Brown, Hugo Romat, Natalie Riche, Samuel Maina, Shamsi Iqbal, S., Siân Lindley, Stephanie Nyairo, Su Lin Blodgett, Sumit Gulwani, Sunayana Sitaram, and Vu Le. 2024. Microsoft new future of work report 2024. *Microsoft Research Tech Report*, MSR-TR-2024–56. https://aka.ms/nfw2024

Buxton, William. 2019. Ubiety, mobility & transitions: The quest for ecological intelligence. *Proceedings of the 2019 ACM international conference on Interactive surfaces and spaces (ISS '19)*, 1. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3343055.3359724

Chui, Michael, Eric Hazan, Roger Roberts, Alex Singla, Kate Smaje, Alex Sukharevsky, Lareina Yee, and Rodney Zemmel. 2023. *The economic potential of generative AI: The next productivity frontier.* McKinsey & Company, 14 June 2023. https://www.mckinsey.com/capabilities/mckin

sey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier. Accessed 4 December 2025.

Dell'Acqua, Fabrizio, Edward McFowland III, Ethan Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim Lakhani. 2023. Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. *Harvard Business School Working Paper,* No. 24–013. https://doi.org/10.2139/ssrn.4573321

Duan, Wen, Shiwen Zhou, Matthew Scalia, Xiaoyun Yin, Nan Weng, Ruihao Zhang, Guo Freeman, Nathan McNeese, Jamie Gorman, and Micael Tolston. 2024. Understanding the evolvement of trust over time within human-AI teams. *Proceedings of ACM Human-computer interaction 8, CSCW2,* 521. https://doi.org/10.1145/3687060

Duan, Wen, Christopher Flathmann, Nathan McNeese, Matthew Scalia, Ruihao Zhang, Jamie Gorman, Guo Freeman, Shiwen Zhou, Allyson Hauptman, and Xiaoyun Yin. 2025. Trusting autonomous teammates in human-AI teams - a literature review. *Proceedings of the ACM CHI Conference on Human factors in computing systems (CHI '25)*, 1102, 1–23. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3706598.3713527

Ehsan, Upol, Q. Vera Liao, Samir Passi, Mark O. Riedl, and Hal Daumé. 2024. Seamful XAI: Operationalizing seamful design in explainable AI. In *Proceedings of the ACM on Human-computer interaction* 8, No. CSCW1: Article 119. https://doi.org/10.1145/3637396.

Erickson, Thomas, Christne Halverson, Wendy Kellogg, Mark Laff, and Tracee Wolf. 2001. Social translucence designing social infrastructures that make collective activity visible. *Communications of the ACM* 45 (4): 40–44. https://doi.org/10.1145/505248.505270.

Erickson, Thomas, and Susan Herring. 2004. Persistent conversation: A dialog between research and design. *Proceedings of 37th annual Hawaii international conference on System sciences*. Big Island, HI, USA. https://doi.org/10.1109/HICSS.2004.1265280

Erickson, Thomas, David Smith, Wendy Kellogg, Mark Laff, John Richards, and Erin Bradner. 1999. Socially translucent systems: Social proxies, persistent conversation, and the design of "Babble". *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI '99)*, 72–79. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/302979.302997

Flathmann, Christopher, Wen Duan, Nathan McNeese, Allyson Hauptman, and Ruihao Zhang. 2024. Empirically understanding the potential impacts and process of social influence in human-AI teams. *Proceedings of ACM Human-computer interaction 8, CSCW1,* 49. https://doi.org/10.1145/3637326

Gao, Jie. Yuchen Guo, Gionnieve Lim, Tianqin Zhang, Zheng Zhang, Toby Jia-Jun Li, and Simon-Perrault. 2024. CollabCoder: A lower-barrier, rigorous workflow for inductive collaborative qualitative analysis with large language models. *Proceedings of the 2024 CHI conference on Human factors in computing systems (CHI '24)*, 11: 1–29. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3613904.3642002

Gaver, William, Thomas Moran, Alan MacLean, Lannart Lövstrand, Paul Dourish, Kathleen Carter, and William Buxton. 1992. Realizing a video environment: EuroPARC's RAVE system. *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI '92)*, 27–35. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/142750.142754

Gray, Mary, and Siddharth S. Suri. 2019. *Ghost work: How to stop silicon valley from building a new global underclass*. Boston: Houghton Mifflin Harcourt.

Guo, Jiajing, Vikram Mohanty, Jorge Piazentin Ono, Hontago Hao, Liang Gou, and Liu Ren. 2024. Investigating interaction modes and user agency in human-LLM collaboration for domain-specific data analysis. *Extended Abstracts of the ACM CHI conference on Human factors in*

*computing systems (CHI EA '24)*, 203: 1–9. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3613905.3651042

Hasker, Steve. 2025. *Future of professionals report*. Thomson reuters. https://www.thomsonreuters.com/content/dam/ewp-m/documents/thomsonreuters/en/pdf/reports/future-of-professionals-report-2025.pdf

Hauptman, Allyson, Wen Duan, and Nathan McNeese. 2022. The components of trust for collaborating with AI colleagues. *Companion publication of the 2022 conference on Computer supported cooperative work and social computing (CSCW'22 Companion)*, 72–75. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3500868.3559450

He, Jessica, Stephanie Houde, Gabriel Gonzalez, Dario Andrés Silva Moran, Steven Ross, Michael Muller, and Justin Weisz. 2024. AI and the future of collaborative work: Group ideation with an LLM in a virtual canvas. *Proceedings of the 3rd annual meeting of the symposium on Human-computer interaction for work (CHIWORK '24)*, 9; 1–14. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3663384.3663398

Hughes, John, David Randall, and Dan Shapiro. 1992. Faltering from ethnography to design. *Proceedings of the ACM conference on Computer supported cooperative work (CSCW '92)*, 115–122. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/143457.143469

Kaye, Joseph 'Jofish', Mariah Levitt, Jeffrey Nevins, Jessica Golden, and Vanessa Schmidt. 2005. Communicating intimacy one bit at a time. *Extended abstracts of the ACM conference on Human factors in computing systems*, 1529–1532. ACM, New York, NY, USA. https://doi.org/10.1145/1056808.1056958

Kaye, Joseph 'Jofish'. 2006. I just clicked to say I love you: Rich evaluations of minimal communication. *Extended abstracts on ACM conference on human factors in computing systems (CHI EA '06)*, 363–368. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/1125451.1125530

Lauro, Quentin Romero, Jeffrey Bigham, and Yasmin Kotturi. 2024. Exploring the role of social support when integrating generative AI in small business workflows. *Companion publication of the 2024 conference on Computer supported cooperative work and social computing (CSCW Companion '24)*, 485–492. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3678884.3681895

Lee, Hao-Ping, 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. *Proceedings of the ACM conference on Human factors in computing systems (CHI '25)*, 1121: 1–22. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3706598.3713778

Li, Chenglong, Xueman Li, and Ting Long. 2025. A systematic literature review on human-AI collaboration in the information systems research. *E-Business. Generative artificial intelligence and management transformation. WHICEB 2025. Lecture notes in business information processingness Information Processing,* eds. Tu, Yiliu Paul and Maomao Chi, 549. Springer, Cham. https://doi.org/10.1007/978-3-031-94184-9_6

Lindley, Siân, and Denise Wilkin. 2023. Building knowledge through action: Considerations for machine learning in the workplace. *ACM Transactions on Computer-Human Interaction* 30 (5): 72. https://doi.org/10.1145/3584947.

Moore, Phoebe. 2018. *The quantified self in precarity: Work, technology and what counts*. London: Routledge.

Muller, Michael, Jessica He, and Justin Weisz. 2024. Workplace everyday-creativity through a highly-conversational UI to large language models. *Journal of ACM* 37 (4): 111.

Nardi, Bonnie, Steve Whittaker, and Erin Bradner. 2000. Interaction and outeraction: Instant messaging in action. *Proceedings of the ACM conference on computer supported cooperative work (CSCW '00)*, 79–88. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/358916.358975

Nonaka, Ikujiro, and Hirotaka Takeuchi. 1995. *The knowledge creating company: How Japanese companies create dynamics of innovation*. Oxford: Oxford University Press.

Nonaka, Ikujiro, and Ryoko Toyama. 2015. The knowledge creating theory revisited: Knowledge creation as a syntheizing process. In *The essentials of knowledge management*, ed. John Edwards. Basingstoke: Palgrave Macmillan.

Nonaka, Ikujiro. 1994. A dynamic theory of organisational knowledge creation. *Organisation Science* 5(1): 14–37. https://www.jstor.org/stable/2635068

O'Hara, Kenton, Michael Massimi, Richard Harper, Simion Rubens, and Jessica Morris. 2014. Everyday dwelling with whatsapp. *Proceedings of the ACM conference on Computer-supported cooperative work (CSCW '14),* 1131–1143. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/2531602.2531679

Orlikowski, Wanda. 2002. Knowing in practice: Enacting a collective capability in distributed organizing. *Organization Science* 13 (3): 249–273. https://doi.org/10.1287/orsc.13.3.249.2776.

Orlikowski, Wanda. 2006. Material knowing: The scaffolding of human knowledgeability. *European Journal of Information Systems* 15: 460–466. https://doi.org/10.1057/palgrave.ejis.3000639.

Russo, Daniel. 2024. Navigating the complexity of generative AI adoption in software engineering. *ACM Transactions on Software Engineering Methodology* 33 (5): 135. https://doi.org/10.1145/3652154.

Sako, Mari. 2024. How generative AI fits into knowledge work. *Communications of the ACM* 67 (4): 20–22. https://doi.org/10.1145/3638567.

Schmidt, Kjeld, and Liam Bannon. 1992. Taking CSCW seriously. *Computer Supported Cooperative Work (CSCW)* 1:7–40. https://doi.org/10.1007/BF00752449.

Schmutz, Jan, Neal Outland, Sophie Kerstan, Eleni Georganta, and Anna-Sophie. Ulfert. 2024. AI-teaming: Redefining collaboration in the digital era. *Current Opinion in Psychology* 58: 101837. https://doi.org/10.1016/j.copsyc.2024.101837.

Seely Brown, John, and Paul Duguid. 2000. *The social life of information*. Boston: Harvard Business Review Press.

Sellen, Abigail, and Richard Harper. 2003. *The myth of the paperless office*. Cambridge, MA, USA: MIT Press.

Shin, Joon Gi, Janin Koch, Andrés Lucero, Peter Dalsgaard, and Wendy E. Mackay. 2023. Integrating AI in human-human collaborative ideation. *Extended abstracts of the ACM CHI conference on Human factors in computing systems (CHI EA '23)*, 355: 1–5. *Association for Computing Machinery, New York, NY, USA*. https://doi.org/10.1145/3544549.3573802.

Simkute, Auste, Lev Tankelevitch, Viktor Kewenig, Ava 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* 41 (5): 2898–2919. https://doi.org/10.1080/10447318.2024.2405782.

Star, Susan Leigh, and Anselm Strauss. 1999. Layers of silence, arenas of voice: The ecology of visible and invisible work. *Computer Supported Cooperative Work (CSCW)* 8 (1–2): 9–30. https://doi.org/10.1023/A:1008651105359.

Strauss, Anselm. 1985. *The social organization of medical work*. Chicago: University of Chicago Press.

Strauss, A. 1988. The articulation of project work: An organizational process. *The Sociological Quarterly* 29 (2): 163–178. https://doi.org/10.1111/j.1533-8525.1988.tb01249.x.

Suchman, Lucy. 1994. Supporting articulation work. *Proceedings of the 5th international conference on Women, work and computerization: Breaking old boundaries, building new forms*. Manchester, UK. https://doi.org/10.1016/B978-0-12-415040-9.50118-4

Sun, Yuan, Eunchae Jang, Fenglong Ma, and Ting Wang. 2024. Generative AI in the wild: Prospects, challenges, and strategies. *Proceedings of the ACM CHI conference on Human factors in computing systems (CHI '24)*, 747: 1–16. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3613904.3642160

Vanukuru, Rishi, Payod Panda, Xinyue Chen, Ava Scott, Lev Tankelevitch, Sean Rintel. 2025. Designing interfaces that support temporal work across meetings with generative AI. *Proceedings of the 2025 ACM designing interactive systems conference (DIS '25)*, 3600–3620. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3715336.3735833

Verheijden, Mathias, and Mathias M. Funk. 2023. Collaborative diffusion: Boosting designerly co-creation with generative AI. *Extended abstracts of the ACM CHI conference on Human factors in computing systems (CHI EA '23)*, 73: 1–8. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3544549.3585680.

von Uexküll, Jakob. 2010. *A foray into the worlds of animals and humans: With a theory of meaning*. Minneapolis: University of Minnesota Press.

Wagman, Kelly, Matthew Dearing, and Marshini Chetty. 2025. Generative AI uses and risks for knowledge workers in a science organization. *Proceedings of the ACM CHI conference on Human factors in computing systems (CHI '25)*, 1199: 1–17. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3706598.3713827

Wang, Ruotong, Lin Qiu, Justin Cranshaw, and Amy Zhang. 2024. Meeting bridges: Designing information artifacts that bridge from synchronous meetings to asynchronous collaboration. *Proceedings of ACM human-computer interaction 8, CSCW1* 35. https://doi.org/10.1145/3637312

Wolf, Christine, & Jeanette Blomberg. 2019. Evaluating the promise of human-algorithm collaborations in everyday work practices. *Proceedings of the ACM human-computer interaction 3, CSCW,* 143. https://doi.org/10.1145/3359245

Woodruff, Allison, Renee Shelby, Patrick Gage Kelley, Steven Rousso-Schindler, Jamila Smith-Loud, J., and Lauren Wilcox. 2024. How knowledge workers think generative AI will (not) transform their industries. *Proceedings of the ACM CHI conference on Human factors in computing systems (CHI '24)*, 641: 1–26. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3613904.3642700

Yun, Bhada, Dana Feng, Ace Chen, Afshin Nikzad, and Niloufar Salehi. 2025. Generative AI in knowledge work: Design implications for data navigation and decision-making. In *Proceedings of the 2025 CHI conference on Human factors in computing systems (CHI '25)*, 634: 1–19. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3706598.3713337

Zhang, Rui, Wen Duan, Christopher Flathmann, Nathan McNeese, Guo Freeman, and Alyssa Williams. 2023. Investigating AI teammate communication strategies and their impact in human-AI teams for effective teamwork. In *Proceedings of ACM Human-Computer Interaction 7, CSCW2* 281. https://doi.org/10.1145/3610072