

Understanding the impact of an organisational LLM agent on the collaborative practices of knowledge work

Kenton O'Hara

kenton.ohara@bristol.ac.uk

University of Bristol

Michael Massimi

Slack Incorporated (United States)

Miftah Khan

Plative

Research Article

Keywords: Generative AI, Large Language Models, Knowledge Work, AI Agents, Computer-Supported Cooperative Work, Work Practices, Enterprise AI

Posted Date: October 7th, 2025

DOI: <https://doi.org/10.21203/rs.3.rs-7601582/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

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 the how these technologies impact on 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 organizations 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 in particular how this shapes collaboration in knowledge work.

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, continues to dominate contemporary discourse about knowledge work futures, as industries seek to understand and adapt to their transformative impact [11, 9, 10, 22]. 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 [11, 38, 39, 40, 46, 51, 52, 55, 58, 59]. These studies offer us initial 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 Chat GPT, and in turn has largely been driven by individual knowledge workers seeking efficiency and productivity gains in their own work. Given such a focus, several other facets of the adoption and use of these technologies have been less explored yet constitute a significant part of how we will need to understand these technologies in current and future 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 and collaborative concern [25, 33, 34, 36, 41, 43, 44]. More specifically, these traditions have done much to highlight the situated nature of these collaborative knowledge work practices and the ways in which they are played out through various configurations of technologies, infrastructures and artefacts [36, 38, 50, 47]. 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 made similar arguments with respect to how we understand new algorithmic technologies in contemporary knowledge work urging closer

attention to be paid to the emerging collaborative and situated practices that surround their use [30, 57]. While there have been efforts to study collaborative experiences with LLMs, these have often typically been focused on task-specific investigations such as, for example, collaborative brainstorming or co-creation activities [20, 24, 31, 45, 53]. These help understand some of the dynamics of collaboration with the presence of LLMs but are removed from everyday situated contexts of collaborative knowledge work. In this regard there is a need to further build on these initial efforts to augment our understanding of collaboration with real world adoption on LLM in collaborative knowledge workflows.

Secondly, while individuals have been able to quickly adopt these tools for their individual productivity needs, at the organisational level, a more cautious approach has been needed for the official introduction of these tools more into organisational practice. With the general off-the-shelf LLM tools, there were inevitable concerns with commercially sensitive data being released into the corpus of training data, hallucination of false information and lack of specificity of these tools with respect to organisationally specific knowledge (e.g. SOPs, policies, customer and project information etc.) [28, 55]. To mitigate these concerns, organisations have begun to develop their own organisationally specific agents that are powered by the same LLMs underpinning the publicly available versions, but grounded and bounded within the documents, systems and data comprising their specific organisational knowledge bases. Given their relatively recent introduction, there has been little opportunity, to date, to study these systems and understand their impacts on the collaborative organisation of knowledge work. A recent study by Wagman, et al [55] offers some exception here, being one of the first to explore a more organisational-level deployment of a bespoke LLM conversation agent developed for a large Science organisation. The work offers useful insight into the ways such an agent is used by science and operations teams in classic “copilot” mode as well as their desire for richer use in science workflow automation. However, while they point to the significant network effects of potential organisation wide deployments, their work does not specifically address the collaborative nature of knowledge work and the impact of their system on those practices.

In this paper then, we present a real-world study of an organisationally bespoke LLM agent developed and internally deployed by an international IT services consultancy. Accessible via Slack, the LLM agent is deeply embedded within the existing organisational structures for project-based communication and collaboration, and interfacing with key enterprise productivity and management systems (e.g. project management, customer relations management and employee relations etc). The study explores how the knowledge workers within the organisation, adopt and use this system and how they situate it in the everyday trajectories and collaborative practices of their knowledge work. In this regard, our concerns are not so much with the notion of AI as a collaborating partner, but rather in its agency in the evolving collaborative practices of the organisation.

2 RELATED WORK

To help situate our research, we present an overview of related work that both supports the motivation for the current research and offers foundational insights to inform our work. While we acknowledge and

are deeply informed by a long-standing tradition of workplace studies within CSCW, we focus here on more recent work pertaining to GenAI. We begin with a look at studies of Gen AI in knowledge work followed by a discussion of gen AI in the context of collaboration.

2.1 Gen AI in Knowledge Work

Since their introduction a number of 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 [11, 38, 39, 40, 46, 51, 52, 55, 58, 59]. Together, these bodies of work offer critical insights into the promises and pitfalls of deploying generative AI in knowledge-intensive domains. Notable here is the research of Woodruff et al [58] who presented a qualitative enquiry into knowledge workers across a range of “knowledge” industries. Focusing on initial use of off-the-shelf Gen AI 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. Being an early study in the area, the work was focused on early practices of initial adoption by individuals and their expectations for how any industrial transformation would unfold. As such, it was not positioned to consider the broader organisational level adoption of these tools and their incorporation into the organisational knowledge bases, processes and tools. Likewise, given the concerns of the work were largely with individual level impacts, it didn’t really explore any changing practices of knowledge workplace collaboration.

In a related study, Brachman, et al [3] also explored adoption and desired future use of LLMs by knowledge workers but focused more specifically on employees from a large international technology organisation. They highlighted 4 key areas of LLM usage including in *document creation/ideation*, *information search and analysis*, *advice seeking*, and *task automation*. Building on this, their participants also discussed future value of LLM-infused 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 unable to offer further insights and details into the emergent everyday practices around these use cases. However, as we are now starting to see the rollout of these value points in organisationally tailored LLM assistants, there is an important need to further detail case studies of their impact on collaborative knowledge work practices.

Other work has sought more specifically to understand and evaluate the impact of LLMs on productivity and quality of knowledge work. For example, Brynjolfsson et al. [8] 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 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. Some alignment with these findings was noted in a related study by Dell’Acqua et al. [11] which explored how GPT-4 affected productivity and quality outcomes on realistic knowledge work tasks by consultants from a global consultancy organisation. The authors nuance their analysis through the concept of the “*jagged technological frontier*” suggesting that Gen AI appears to excel at some tasks while failing at others of seemingly similar difficulty. With this in mind, 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 lower 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 [46] and Lee et al [29]. Simutke’s work highlights a potential productivity paradox of Gen AI 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 [29] develop these ideas further, specifically in relation to critical thinking aspects of knowledge work. They introduce the notion of “*AI stewardship*” suggesting a need for support tools, and knowledge worker skill development, to guide and monitor Gen AI to achieve high quality outputs in the contexts of their work. While this research offers valuable insights and implications, much of the analytic focus still remains on 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 [cf., 48, 49] of knowledge work that span across multiple systems, workers and organisational processes. As we shall discuss through our own work, this can have important implications for how we both understand and design for such AI stewardship.

Other more recent research [59] identifies critical challenges faced by knowledge workers when using generative AI for data navigation and decision-making. In particular, it highlights the difficulties 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, we will see in our research how these conclusions offer resonance with some of our own findings of real-world deployment and use.

2.2 Gen AI in Collaboration

In terms of understanding generative AI use as a more collaborative concern, there is much less attention 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* [13, 14, 18, 19, 23, 42, 60]. Much of this research has focused on particular elements of collaboration dynamics between humans and agents such as differences in communication strategies, social influence, trust, interaction and agency. While these social dynamics are important to our overall understanding of AI in team collaboration, their aim is not to elucidate the everyday situated practices of AI use in knowledge work. Other research has focused more on the use of gen AI tools in human collaboration tasks without necessarily invoking the AI as an additional team member [24, 20, 31, 45]. These efforts have tended to focus single bounded synchronous tasks and often within the context of design-oriented activities such as brainstorming [31] ideation [24, 45], co-creation [53] and qualitative analysis [20]. So, while these tasks may be representative of certain knowledge work activities, they are treated as single point isolated events. As they do not situate these activities within the more typical trajectories of interdependent communication, coordination, sense-making and production activities that comprise organisational knowledge work.

Building on the work of Strauss [48, 49], other aspects of CSCW research have adopted a more longitudinal perspective on the nature of collaboration in knowledge work. Strauss introduced the idea of work *trajectories* referring to the temporal and socio-material unfolding of collaborative work. Of significance within these ideas is the additional notion of *articulation work* which relates to 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. These ideas were key to early CSCW work which did much to highlight the emergent, contingent and situated nature of knowledge work and the need for technologies to support the key aspects of *articulation work* along project trajectories [e.g. 41, 50]. Schmidt and Bannon, notably point to the importance of standardisation and structure in organisational processes and artifacts in supporting articulation work along collaboration trajectories – something that has implications for understanding our own work.

More recently, these core concepts and ideas have been revisited in the context of emerging Gen AI capabilities. Wang et al [56], for example, introduce 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, the work highlighted challenges

with information capture during meetings and impact on engagement, and difficulties re-establishing context or finding desired information. Based on their analysis, they established a number of key design principles for meeting bridge systems, such as the incorporation of different linked data types, structure to help users explore information, the ability to share customized information and artifacts, connection to contextual information outside the meetings, and the ability to evolve information after meetings. With these principles in mind, the authors discuss how emerging LLM capabilities could be used to address these challenges and principles. While the work doesn't yet follow through on these arguments to develop technologies and to understand their use in practice, as we shall see, they nevertheless provide some useful foundations for our own work.

In related work, Asthana et al [1] seek to further explore these ideas through the development of 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 highlight the importance of different but complementary recap levels with *Highlights* and *Hierarchical* models that can be better adapted to the different contexts of subsequent work. They emphasise the importance of the recap document as a natural artefact for collaboration and consequently the importance of seamless integration with other everyday corporate tools of knowledge work. With this in mind, the authors suggest significant potential of these recaps in transforming working practices of knowledge workers particularly with integration into corporate collaboration tools such as Microsoft Teams or Slack. The authors acknowledge that understanding these practices remains an important opportunity for future research – a call that resonates strongly with the motivations of the work presented in the current paper.

Further recent work by Vanukuru et al [52] has sought to build on these ideas to more explicitly examine the temporal nature of knowledge work that connects together meetings and projects across time. Drawing on core concepts of *retrospections* (looking back) and *prospections* (looking forward), the research presents a series of Gen AI enabled concepts (*Instant Recaps*, *Adaptive Meeting Handoff* and *Project Browsers*). These concepts are designed to collectively support a chain of intentionality and reflection across meetings, where summaries from one meeting 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 provide intriguing inspirations for understanding and coherently supporting temporal features of knowledge work with GenAI, the opportunity to understand their impact on every day organisational practices remains something for future work.

Approaching these issues from a different perspective, Lindley and Wilkins [30] explored opportunities and considerations for the adoption of machine learning and Gen AI in knowledge-based organisations. Their particular concerns were with the use of AI to generate organisational knowledge bases drawing on the actions and practices comprising everyday knowledge. Furthermore, they looked at how such AI infused knowledge bases could enable the development of new practices that enhance how knowledge work is accomplished. Grounding their arguments in longstanding discussions of organisational knowledge and knowledge work [e.g. 33, 34, 35, 36, 38, 44], the research frames knowledge work in

terms of *actions* and not just in terms of outputs and documents. Notable here was their consideration of both *Collaborative Actions* (CAs) and *Knowledge Actions* (KAs) as potential signals for building ML knowledge bases. These included things such as seeking help, gauging expectations, approvals, planning and strategy work, coordination and articulation work [cf. 41, 48, 49, 50] required to stay in sync, generate and capture information and transfer information. The significance of highlighting these actions is an acknowledgement that much of organisational knowledge is tacit, embodied and socially enacted. With this in mind, the argument is that machine learning systems can learn both from explicit inputs comprising the organisational knowledge base as well as more implicit signals found in the socially enacted practices of their production and use. Furthermore, they argue that machine learning tools in the context of knowledge work, should be considered not just as something to work with but as something to work through [6] emphasising the need to consider how design and adoption of ML and GenAI tools transform and support the *processes* of doing knowledge work. With this in mind there is a strong need for us to develop a deeper understanding of how this is happening in real world organisations in the context of their everyday collaborative workflows, not just their individual workers

Adoption at a more organisational level has been somewhat more cautious with concerns of data privacy and security as well as with output specificity. However, such concerns are being increasingly addressed with the development of bespoke AI agents, powered by these more general models but tailored and grounded in organisational knowledge bases. Research into the adoption and use of these tools, though, remains limited [e.g. 55]. Wagman et al. undertook a study of an LLM powered chatbot, Argo, that was deployed in the National Science Lab in the US. The research combined a survey and interviews to capture initial perspectives on the adoption, uses and concerns relating to the tool as well as insights into envisioned future use. The research highlights the importance of its use for writing tasks, aligning with previous research. Of note though is the particular need for working with various structured forms of documentation and reporting that are bound to particular processes of scientific practice. This has significance for our own findings and the adaptation of processes and practices related to structured reporting. Likewise, the study highlighted a need to be able to mine and extract insights across large unstructured sets of documents. This included scientific literature and public data sources but also an expressed desire to extract insights from meeting transcripts and interrogate organisational procedures, policies and documents. While the research begins to point to some interesting current and envisioned use cases for an organisationally tailored level generative AI assistant in a specific science organisation, there is not a detailed elucidation of the everyday collaborative work practices that have evolved through its introduction. There remains an important need to build in this work and further understand the emergent work practices of specific organisations adopting more organisationally tailored Generative AI Assistants.

The work of Wolf and Blomberg [57] offers some further insights here. While their research predates the more recent introduction of LLMs, it nevertheless focused on the development of an NLP (natural language processing) infused tool in an IT service provider organisation with particular reference to the IT services procurement process. Such a domain is of particular relevance to the concerns of the current paper that is also grounded in the collaborative knowledge work of an IT service provider organisation.

The work is of further significance in its analytic orientation. Drawing on foundational studies and approaches in CSCW, Wolf and Blomberg, highlight the need to attend to “the unfolding and tangled relations between situated work practices and the emergent technologies that we might design to support, augment, or transform such practices.” In this respect their research offers a rich description of these situated practices, engaging with the workers to understand how they consider the appropriate application of the developing NLP tool in relation to their specific everyday work practices and tasks. 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 are leveraged and transformed through human interaction. Secondly, in the sense that the IT solutioning work was seen as a fundamentally “team sport” that would require algorithmic outputs to be integrated in ways that enabled their collective interpretation and transformation through the collaborative practices of the team.

Building on this prior work we present a 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 work practices. The study explores how the agent is implicated in an evolving set of collaborative practices of the knowledge workers and how these play out across the temporal trajectories of their work. We begin with an overview of the organisation and the LLM based agent, *Plaito*, that they have introduced to their working practices. We then present the findings arising from the study and discuss their implications for how we need to think about and understand these systems in the collaborative organisation of 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 wide variety of clients with specialist expertise in financial services, professional services, nonprofit organizations, technology, media, and consumer sectors. They have a global presence with offices and clients distributed across 5 continents.

Within the organisation there are 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. With this in mind, the employees rely heavily on *Slack* and *Google Workspace* (Meet, Calendar, Drive, Docs etc). 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.

3.2 “Plaito” LLM Bot

Plaito is a customised AI bot developed by the organisation. It is based on Open AI but has been 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 accessible from within Slack, where it can be integrated in the communication and collaboration channels that underpin significant aspects of 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 content of these messages can include URLs, links to documents or specific content linked within the enterprise software tools.

Over and above the natural language prompt input box, Plaito also has a number of “short-cut” buttons which provide quick standard access for frequently acquired information from Confluence, Salesforce, Netsuite, Tableau, LinkedIn and Jira. 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 structure 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 prompt-based interrogation.

3.3 Interviews

To orient the research team, a senior leader in Plative provided us with an introductory overview of the organisation, its work, and the key stakeholder roles involved. Volunteers from these different roles were invited to participate in the research. In total, fourteen participants were recruited to be interviewed (see Table 1). An initial interview (approx. 30 mins) was conducted with all fourteen participants. During these interviews, we discussed their role at Plative, their everyday activities and tasks, the key tools and artefacts of their work and how Plaito more specifically was integrated into their workplace collaboration. Following this interview, participants were, availability permitting, invited to participate in a second interview 2 weeks later – nine were available to continue further. Over the course of those two weeks, the participants were asked to collect 3-5 concrete examples of using Plaito in their everyday work, documenting their interactions with screenshots. These concrete examples formed the basis of discussion during the second interview (approx. 60 mins). Grounding the discussion in these concrete examples, enabled participants to recount the details of these episodes, the context of the work, its

tasks, any unfolding communications and collaborations with team members or clients, and the situated use of AI in this context. Where appropriate, these conversations were supported with illustrative screen shares. Interviewers would prompt participants for additional details where required as well as explanatory accounts of how and why things unfolded in particular ways again with reference to people, artefacts, and organisational concerns. All interviews were conducted online using Microsoft Teams or Google Meet. Automated transcripts were produced using inbuilt transcription software, and subsequently reviewed, corrected and anonymised in preparation for analysis.

Table 1. Participant Roles and Interviews

Participant	Role	1 st interview	2 nd 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

3.4 Analysis

The interview transcripts were analysed using a reflexive thematic analysis (RTA) approach [7] that acknowledges the researcher’s active interpretative role as fundamental to the meaning making and analytic account produced. The approach is inductive with constructionist concerns whereby meaningfulness and significance in the data was given extra weight with recurrence of issues across different participants. Likewise, while our approach primarily reflects the experiential perspectives and understandings of our participants in the collaborative organisation of their work, there is an additional concern with reimagined future configurations of their collaborative work that might be enabled by

further machine learning innovation. Following the key stages of the RTA process, we began our analysis with a data familiarisation phase, 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, labelling relevant excerpts in the data with initial codes. These codes were iterated on and clustered together to form into high level themes.

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 are able to introduce their own workflow and process improvements by generating “documents as meetings”. Finally, we discuss news 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, pretty much all meetings are conducted online via Google Meet software. By default, the culture within the organisation is that all meetings are recorded and automatically transcribed with LLM-enabled transcription software. From this transcription document, Plaito generates another document that provides a structured summary of the meeting.

These documents are attached to the associated Google Calendar event as well as stored in the Organisation’s Google Drive knowledge base. As we see in Figure 1, the meeting summary generated by Plaito consists of several standard sections that have been deemed relevant by the Organisation. These sections include *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 are populated by a bulleted list of the key points. There is also a highlighted disclaimer that this is an AI generated summary and that results should not be distributed beyond the immediate attendee list or to external persons prior to a review and accuracy check.

Through this process of transcription and summarisation, the meeting is transformed from an ephemeral event into a persistent artefact – a working document that gives the meeting material form^[1]. This material transformation of meeting to document format has a number of significant effects on how collaborative elements of the work are organised. Documentation remains a significant part of how the organisation demonstrates its professionalism to its clients. In the first instance, this material transformation of a meeting to a document rather straightforwardly alleviates the burden of having to take notes during the meeting. As a consequence of alleviating that burden, participants reported that they were able to be more “*invested in the call*” and more actively listen to the concerns and requirements of the client.

“So historically I would be on that call frantically taking notes and as you know 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

Importantly here, this was not simply about the pragmatic and cognitive concerns of 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.

“100% yeah, like when I trained my teams like one of the, I remember one of the sessions that I had at NetSuite was I had a Rep that had a really high close rate and the reason that that Rep was so good as I kind of dove in 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, you know what I mean? Like notes can be distracting. He was really good at doing both.” 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 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 utterances they captured. 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 like the goal like OK, like sometimes we'll jump into a meeting, forget to press record and then we'll press record and we'll go and like introduce the meeting as if like a new stakeholder just joined right that like, oh, we're talking about this and blah blah blah and then we'll get it back into our discussion. And then I do find times where we're like they're kind of like crosstalk, 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 like say the ticket number. I'm like, great, OK,

it's there. So yeah, I'll find myself talking, I guess. Yeah, talking to the AI that's going to process the transcript.” P5 Delivery Manager

The quote highlights how talk within these calls isn't just reflexively accountable to colleagues and clients in the meeting, but also to the interpretive procedures that will ultimately render it legible. This “legibility” is more than 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 display and sustain the orderly properties of this new meeting setting - namely, that speaking is formatted to pass as properly machine-readable talk.

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

“So, I've noticed before, so I'll take notes on my own side. I know that *Plaito* will send me the summary stuff, but sometimes they've missed things like they can't capture everything. So, there's really small things and I make sure that I like write it down on a separate screen. So, I don't forget like if a client has a question and I'll say yeah, I'll send you some documentation. And then I don't ever bring it up again. *Plaito* sometimes doesn't catch that. And then I don't want to say that I'm going to do something and not follow through on it. So, I make sure that I have like a little list. Going on my other screen and then at the end of the call I'll say OK, so to summarise, I'm going to do this. I'm going to do this, I'm going to do this. You are going to do this, you are going to do this. So just so that it reads the transcript or I don't know how it exactly it comes up with all of this. If it reads it or if it listens to it, I don't know. But I make sure that my action items I say multiple times throughout the call...I try to just continue to repeat the important things. So, I know for a fact at all be caught on the summary.” P7 Business Analyst

There is an interesting sense in which performatively clarifying and structuring content to render it meaningful to the AI, also served the purpose of rendering things salient to others within the meeting signalling their significance to the work.

We can extend these arguments further to consider how these assemblages of human-AI *intra-actions*^[2] deeply reconfigure the collaborative trajectories of organisational decision making and practice. Again, these reconfigurations are not simply passively following an arrow of time in which a meeting pre-exists the summarisation and the summarisation just acts on the meeting. Rather, it is in the possibility for summarisation and opportunity for agentic interrogation of the summary that the whole collaborative trajectory is transformed. To illustrate, let us consider an example articulated by P11, a talent acquisition specialist. In this scenario, P11 described an episode involving a screening interview for a job candidate.

Her task was to provide an appropriately structured account of the candidate that could subsequently be reviewed and acted upon 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 who had previously agreed with her a list of questions and information requirements they wanted to know. While she was accountable to the hiring manager's agreed list of questions, she recounted how the conversation with the candidate no longer needed to be so rigidly structured by questions. In the knowledge that she would be able to restructure things later with *Plaito*, she explained how she no longer needed to worry about deviations from the structured conversation during the interview. In this respect she was able to adapt the conversation fluidly and attentively to the candidate.

"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 one second here one 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 instead of having it pointed out for the answer. That makes the questions into my conversation instead of having it ... Some of these questions I wouldn't even ask to the candidate. I could actually kind of work it into a conversation where they kind of don't know the questions being asked and they'll be able to answer it... I think so I can actually listen to...what they're saying and have a little bit more deeper concentration with the conversation with this person instead of kind of having to type something out about an answer they said or" - 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. This 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 helps 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 I don’t have to go back to look at the answers or find a specific word in the document. So, I know that AI can pull that a lot more faster. which saves me time which means that 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 kind of thing. So that takes the summaries I can do within whatever half an hour an hour depending how many interviews I do a day, but it saves me a lot of time.” P11 TAS

P11’s expression of being able to “depend” on the AI here highlights how it became entangled in the material-discursive intra-actions and practices of determining *relevant* candidate competencies and ultimately “*fit*”. It doesn’t simply follow on from the candidate interview as an extractive function but can be seen to influence the way the candidate interview is conducted and what is made relevant to the hiring manager in their ultimate decision making.

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 than simply retrieving information. 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. To illustrate this, consider an example recounted by one of the directors. In this example, the participant discussed a core activity that they are required to do as part of winning new client business, where they need to respond for a Request for Proposal (RFP)^[3].

“Typically, however, people are hamstrung by the capacity of technical resource and so what this [Plaito] has allowed for is to again be more efficient with time. Where, let’s say historically it takes six hours to review and respond to an RFP because of all of the details and all of the stuff that’s there. And because

of that, because of the different Subject Matter Experts (SMEs) we now are looking at the six hours by 4 folks across the team that's also pulling them away from other activities, other [work] cycles that 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? No problem. So, I now have that Bob has spent 5 minutes giving me that list as opposed to two hours going through the document looking for that information. Myself as a client partner, I've been empowered to now 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* if you will, around the critical information so that you can get what you need. So, I think everyone has gotten time back in their day so that they could be more impactful than some of the redundant tasks that we do, we've just kind of elevated where we spend our time more often than not across the entire client acquisition to project delivery.” P9 Senior Director

What we see articulated in this quote is how, prior to the introduction of Plaito, the process of making sense of the RFP and preparing questions and responses for the client would typically have taken significant amounts of time involving several SMEs inspecting documents and being on long calls with each other and clients for which they may only be required for small amounts of time. Underpinning this more traditional response process was the presumption of co-presence and synchronous availability of SMEs within meetings to be able to meet the tight turnaround times of these responses. But with the introduction of *Plaito*, P9 describes how these presumptions of co-present and synchronous participation have shifted and where new opportunities for asynchronous and remote collaboration reduced time demands on SMEs and the overall coordination burden. This in turn allows for a much faster turn around time in the response to the RFP. Specifically, one of the SMEs, Bob, was interviewed by P9 for 5 minutes to extract their priorities that they needed to address. P9 was then able to use Plaito to interrogate the massive RFP document with prompts derived from the interview responses from Bob. Furthermore, in subsequent calls with the client, he no longer required that Bob be there in person. Rather, in those calls, he was able to leave what he referred to as “*breadcrumbs*” in the meeting that would enable Bob to work with Plaito to interrogate the meeting transcripts jumping directly to where his prioritised information needs could be validated – asynchronously and remotely. Here the ability for both P9 and Bob to interact with Plaito to interrogate documents, recordings and transcripts reconfigured the very ontology of expert participation in work of responding appropriately to the RFP.

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*. To illustrate, consider the following example, in which P11 has been asked

by one of the senior directors to create some new Linked-in messages to contact potential new talent for the organisation. The scenario begins with a call between P11 and the Senior Director [Jack]. During the initial call to discuss this messaging, the Senior Director had outlined several content and style suggestions that the new messages should incorporate. With this call having been recorded, P11 engaged in an exchange with *Plaito*. In the first exchange, P11 asked Plaito to summarise what the Senior Director had articulated during the meeting about the LinkedIn messaging aims, content and style:

P11: “Summarise for what [Jack] is suggesting on the messaging we should use for Linked-in [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 continues for a couple more prompted refinements. At this point, P11 copies 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 responds 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 minutes on your calendar to discuss this opportunity?”
- “Are you open to a brief conversation?”

No need to add “No pressure, just checking” etc.”

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

What we see in this scenario is how the ontology of writing has been fundamentally transformed by the coming together of human and AI agencies in the collaborative production of content. Notable here 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 material manifestation of the meeting, P11 has engineered a form that can then be invoked in a follow-up prompt where it is used as a foundation on which the prompted content generation can be constructed. Here, her prompt neatly encompasses the collective agency of P11, Director and Plaito. The follow up prompt suggests that the initial response by Plaito, doesn’t sufficiently capture the voice and suggestions and that it should expand more with what “[Jack] is suggesting.”

Once P11 is satisfied with the output suggestions, she brings back in the voice of the Director, by sharing the content in a new Slack chat with him, acknowledging her work with Plaito and inviting his thoughts. His positive response in the chat approves the content with the need for only a small addition to make it suitable for posting. In this respect, then, the finally produced content was not simply authored but rather collectively enacted. Each meeting, message, comment, and prompt were shaping what the content would become — foregrounding certain rhetorical styles, prioritising certain information and 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 person, tool or interaction but rather distributed and emerging through a 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 along project trajectories, significantly reducing the burden of ongoing project documentation. A key example of this was seen in the process of translating requirements from a client *Discovery* meeting into trackable project management objects in Jira^[4]. All projects within the organisation are initiated by a discovery phase in which they engage the client in one or more meetings to understand their specific requirements. Based on these discovery meetings, the requirements need to be formalised into a series of *User Stories* with associated “*Solution components*”, “*Acceptance Criteria*” and “*Level of Effort*” required – facilitating client engagement, resource planning and project execution. These components have to be initially documented in a project *user stories spreadsheet* (see Figure 2). The spreadsheet is a collaborative document that is authored by multiple Business Analysts and Project Leads across different project workstreams. Once it is completed, it is then submitted as an Agile request as part of a larger workflow in which the spreadsheet entries are translated 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.

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

“I used to be a BA before getting to being a principal now and being a BA my nightmare was writing user stories because there were just so many of them. and then trying to get the syntax right, trying to think about all the possible acceptance criteria that can ever exist, you have to do that. But with Plaito right now, we're able to cut documentation time in half and get that out of the way.” P1 Principal Consultant

P1 describes 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 build the *Acceptance Criteria* and then add the *Solution Components* and the *Level of Effort* for each item in this sheet. This is based off of the statement of work and based off of the discoveries that we've had with the customer... So, when Plaito writes those user stories, we can update the spreadsheet based off of 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 so it will take a couple of seconds to go try to get me some answers in Salesforce. And once it's done, it will come back and it'll 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 that it all 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 are 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. So 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 what the priority of this is, 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

Important in these sequences is how the collaborative knowledge workflow remains mediated through an established set of coordination artefacts and steps. The spreadsheet, for example, coordinates information assembly and documentation across a number of different teams and project leads, as well as collective review and validation before passing on to the next team for subsequent processing. These workflows often reflect hard-won local knowledge and tacit practices built up over time and the existing coordination artefacts (e.g., shared documents, ticketing systems, status boards, informal channels like Slack) serve as scaffolding for mutual awareness, accountability, and temporal ordering of work. This bottom-up “remixing” approach in which Plaito is used to transform information into new workflow-appropriate forms, acknowledges the important roles of these artefacts and remains sensitive to the situated coordination practices and artefacts already in place. It allows users to appropriate the LLM in ways that make sense locally, adapting its outputs to the norms and constraints of their team, role, or domain. It also supports micro-level workflow improvements without requiring global consensus or formal process redesign that has the potential to be misaligned with everyday realities. By integrating LLMs into existing workflows through familiar coordination artefacts, adoption is facilitated and not disrupted. Users can see immediate value and retain agency in shaping how the AI is used enabling for adaptive evolution of workflows. Over time, these micro-adjustments can lead to more significant transformation, but through remixing, they emerge from within the practice, not from outside mandates.

4.5 Documents as meetings and process improvement

Thus far we have seen how the materialisation of meetings through recordings, automated transcription and AI-based interrogation transforms many aspects of the production, sense-making and accountability of collaborative work. As participants have become familiar with these capabilities and incorporated them into their everyday practices, they described additional ways that they have appropriated the tools for collaborative process improvement. A notable example of this was found in their introduction of new templated documentation of repeated organisational processes that were developed to help structure the collaborative work of the meeting with a view to populating these documents. In this example, P6 described how project teams on client projects will have a *Design Review* meeting in which the team discusses automations they have been working on for a client. In this discussion they consider different candidate solutions to address the client's automation needs and articulate their respective pros and cons. To help with these meetings, P6 recounted how she prompted Plaito to create a Design Review template document with appropriate sub-headings designed to effectively capture and structure the outputs from any Design Review meeting.

"So what we've done with it 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**, right. 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, why not. That is something that we want to use. So, this one here is highlighted in red [solution 1] because what we had initially created was time consuming - it wasn't very efficient. It might slow down your apps. And then we had a couple of other solutions that... We essentially put together two different solutions and were able to speak from these interviews with our solution designers and...our developers to then be able to dictate into this document what we're going to do and why it works the best." P6 Consultant

What we see here is how this created template, which might once have been a tool for post-meeting documentation, is now operating as a live, conversational resource during the meeting. The structure of the document and its various subheadings were used as a visual reference during the meeting, organising topic sequencing, determining what should be discussed when, and framing whether talk had achieved a level of adequacy for documentation. Accordingly, the participants in these meetings organised their talk to be hearable by the LLM agent as orderly, complete, and categorisable under prospective headings such as, in this case, "*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, textualisation and documentation. Participants made their contributions not merely to advance substantive discussion during the meeting but to populate an anticipated text – talking the document into existence.

"Yes and it's so fascinating too because when I'm in the meeting I'm like 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] right so we're thinking in the context of actually using it

for specific purpose of using AI and I do talk to myself in the recordings. And I say note-to-self XYZ. And so we've adjusted our thinking based on how we're going to use the AI in the context of the recording because we try to record everything, because it gives us a summary of what our meeting is. But now we'll be able to specifically use a template for the solution design and then be able to lay out everything in that recording. And I like...because if I don't understand exactly what they're talking about, we can still speak to it in the context of the recording....I think and before using templates and recordings like I writing this document would most likely try to understand through and through what our developers are talking about which is a little harder for me because I don't know anything about code. They know very specific things that I don't. So, with using AI like they can tell me and then the AI can pull the exact context of the language” - P6 Consultant

We can see here how participants’ methods show that they are operating under a prospective accountability towards an AI-generated document. Their talk is produced just to make sense in the moment, in the meeting, but must be 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 a 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 AI translation have agency over the meeting talk exerting a structuring force through the practical methods participants use to accountably orient to them.

Others discussed how they were able to create 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 which captured key job profiles within the organisation, 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 process and content.

“So our HR team has put together what we call like our career framework which is, you know, a bunch of bullet points of what we expect in a senior consultant versus a consultant versus a business analyst, right. So they have like these bullet points, right. And so 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, right? And so that has been that that has taken a lot of time off of people's plates, my plate, the hiring managers as well, because at the very least gives us some bare bones to work with and we know for certain that it lines up with what we have identified as our career framework. So what I'll do is I will take our career framework and I'll go to Plaito and say 'based off of these bullet points or based off of these characteristics for senior sales force consultants, please create you know a job description for me'. You

know with these with these different bullet points and it'll just pop something out. You know, a beginning. You know, responsibilities, qualifications, all of that stuff... I will share with them [hiring] managers like what Plaito has popped out. We'll pretty much,.. like the way we do it is we put it in a Google doc and then we just like collaborate together on it and they'll add things or whatever and we'll just like verify and ensure it's not completely out of whack.” P14 Head of Talent Acquisition

The same framework document provides the basis for other aspects of the HR pipeline enabling a consistent presentation throughout. P14 goes on to articulate how the framework is used with Plaito to ground the generation of subsequent interview questions that the hiring manager will be able to ask – in ways that are aligned with the career framework:

“One way that AI has helped in the later interview process is, you know, actually being able to create interview questions that the hiring manager can ask. You know, in the second stage or third stage interviews based off of career framework and different like core competencies that we create people on during their time here at Plative. And so that's been sort of a game changer cause a lot of people don't know how to interview and it's like you know...Yeah, , 'cause this was this was like a big thing for us before, like trying to create like more structured interview questions cause theoretically we should be asking the same questions to every single person to get like an unbiased sort of view on them. And it was a little bit before you know, maybe one person would ask the same questions, but if there's more people in the process they might not be asking all the right questions. And so we kind of went on this journey to create just more structured questions that everyone is asking people. And before Plaito, it was like, you know, I was sitting down, hiring managers, being like, OK, like, give me a list of questions and that was a disaster because they have to take up their time to create these interview questions. Now what I do is I take the job description and I say ‘based off of 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...So it just makes obviously everyone's life a lot easier and it's literally based off of the career framework that we have put into place.” P14 Head of Talent Acquisition

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, are responsible for multiple, projects, resources and client engagements. In the course of their everyday work, the achievement of *situational awareness* across these projects is fundamental to the successful performance of their role. This is variously achieved through a complex weaving together of heterogenous practices to make overall sense of the project and determine action. These can include things such informal talk and chat-based retrievals of status checks, reviewing 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. important, prioritisation and urgency) and anticipating potential future issues before they arise. Situational awareness, in this sense, was not just something attained but something that is an ongoingly and actively achieved practice. The

methods for doing this were bound up in what we might consider 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 fragmentation constrained what was visible, knowable and actionable for them. Their previous methods and practices for producing their situational awareness consequently required them to navigate across these disparate sources and actively stitch things together. This was 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'm also. I also like to stay on top of all the projects that are going on in the other parts of the business, that sort of maybe 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 'cause I've missed some stuff where people tagged me but I have a good overview so I have to oversee so much information. It's all going to be up in the noggin there that..., that recap and the summarisations of channels and stuff that have given me a good..., a way to kind of get through that all quicker... Slack is our front end to query all this information” - P2 Senior Director

The introduction of Plaito within the Slack platform was able to offer 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 this, they were able to engage in new forms of *aggregative surfacing* and *temporal stitching* – whereby Plaito was being used to provide managers with consolidated project-relevant views and summaries across different information systems and across different temporal periods of the projects. In this respect, the manager's situated use of the agent's outputs change the way they are able to sense, sample and perceive their organisational world– to achieve a new *managerial umwelt* where projects, clients, resources, deadlines, risks, and opportunities are stitched together into continuously surfacable, queryable and remixable forms.

“If I want to ask for regarding a project, a specific project then it's I could I could click on status, I can click on risk, budget invoices is 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 - to, you know, if I click on “Status” [button] it's going to do a reply asking Plaito to summarise the project status. After reviewing all project trackers budget, it comes through as this [name] ‘*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]*’ ... Yeah, I'm using just to identify basically what's going on with the project, different projects, et cetera. Yeah, in general, let's say, if I'm going on a vacation for five days, six days, and I want to get a summary of

what happened in the last five days. So I would say, tell me, tell me, Plaito, 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 significance here was the ability to provide some indication and measure of project and meeting sentiment. Based on meeting transcripts, Plaito 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 offers 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 are times also, where Plaito may over-index on particular linguistic features of the conversation that may skew the overall score and summary in both positive and negative directions. While such potential for sentiment misrepresentation might presume the development of an eventual mistrust of the tool this was not the case. Indeed, sentiment analysis occurs openly with full acknowledgement of any potential infallibility, continuing to function as an important device in the hierarchical organisation of client and people relationship management.

The practical reality of this is that senior management in Plative "*cannot be in every place they need to be at any given time*" (P9) to deal with all clients and associated programs of work of different complexities, stages and resource demands simultaneously. As such, they are not always available to participate in every meeting occurring within their client portfolio which can be limiting to the development of their up-to-date situational awareness and consequently their ability to intervene and respond to concerns. What we saw in the study was how Plaito's sentiment analysis and summary offered a way to deal with this. Any call rated 3 or below is automatically flagged to senior management in a private Slack channel. Of significance here is that the alerts and summary are not simply taken at their immediate face value, with staff being literate to the potential for Plaito to sometimes over-index on certain words and phrases. What is important, however, is that these negative flags were demanding of some sort of explanation and it is this which consequently initiated an engagement between the senior manager and their staff (as well as the client if required) to understand the situation further.

"These reports actually go to [my manager]. They will read it and it 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. In the first instance, anticipation of the escalation and enquiries by senior management prompt the consultant to reflect on the meeting and try to understand any situational factors that may have contributed to the score. These may be factors inherent in the client relationship or factors that may explain the AI's flawed assessment^[6]. Through this prompted

reflection, the consultants were able to pre-emptively prepare a response to any queries from leadership that would be likely to arise. Second, seeing the stream of scores appear in the Slack channel provided mutual awareness – all of the leaders could easily see all of the sentiment scores in near-real time. Third, it allowed the appropriate senior leader the opportunity to investigate.

“I'm probably 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, if you want to call it the truth, actually lies, but then also how you get towards a resolution. So, Plaito kind of gives a referee's view of the issue and that's where I used objective versus what client experience or 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 at times before you jump into it. **It also allows me personally to go into a call with context yet still seek to understand.** And with that, as mentioned, many people want to be heard. So, it's like I might know why we're here, but it's given me a tool to actually navigate that conversation with a baseline. And I could do that almost 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 calls and now four hours have gone by before I've addressed this burning issue or maybe I prematurely jumped in and had that conversation but I would have benefited to get more of the context or detail. So it just increases the chance of productive or successful interaction at a critical point. And that's what the signal channel allows us to do.” - P9 Senior Director

The above quote highlights how such conversations can be difficult to navigate but that the sentiment analysis provided “*data to leverage*” (P9) - a grounding on which they could be tactfully and sensitively crafted. It is not just the actual signal that is important here in gracefully managing these conversations. It is also in the presence of the plausible AI fallibility that one can approach these conversations safely and without confrontation because there is a shared understanding that AI flaws (rather than consultant performance) could be a potential explanation for the enquiry.

The utility of these sentiment signals is not simply bound to the labelling of a single meeting. Rather, these are also considered in more longitudinal terms in ways that come to represent the ongoing relationship trajectory with the client. We see this in the following extract in which a senior executive, experiencing an unforeseen problem with a customer, will use Plaito’s “Ask Salesforce” prompt in Slack to provide a sentiment overview of a project over the previous fortnight.

“And I look at my day well in advance of being in the day, but sometimes where there's like something spontaneous that's thrown in. Absolutely, that's kind of a knee jerk reaction now. It's like, OK, who am I talking to now, if it's a current customer, as an example, we have what is our “Ask Salesforce” [prompt button] where we manage all of our project information. So, with that it's capturing financial status of the project risks on all of our calls are recorded and transcribed, and we assign a sentiment to each and every one of those calls. So, in me hitting that button it's gonna give me the sentiment over the last two weeks, so I can see if it's two weeks of green calls, which would be a four or higher. Great. If I see three

Reds, 2 yellows, 1 green and onwards, that's also just gonna be a data point that's gonna allow me to understand what I'm walking into. And then conversationally, I can clarify what were the risks, concerns, issues brought up in this call on October 22nd, so I can go into it with context... So, it will actually take a snapshot of Tableau, which is our BI tool. So, take a snapshot from our sentiment dashboard that we've created, but then it will provide in plain text a synthesis of that information. So, it'll look at that snapshot of the week. It'll talk about the red call the series of yellow calls, or the green calls. So, it'll provide you with context in relation to the calls from that snapshot in time, which we've created the parameters to be. It's either weekly, weekly or biweekly that it's the context window. But yes, that's what it'll do. All of this is in Slack, by the way as well.” - P9 Senior Director

Apparent in the excerpt is how the senior director may be confronted by these issues suddenly out of the blue, with little time to react. From within the Slack communication channel, the senior director was able to quickly prompt Plaito, which grounds its response in the current client-specific channel, providing a longitudinal visualisation and description of sentiment that rapidly contextualises the client issue they have encountered.

A further effect of these sentiment scores being visible by senior management is that they encourage the consultants and business analysts to consciously try and achieve positive sentiment scores through the ways that they conduct themselves on calls with the clients. In this respect, they function as a high-level accountability tool.

“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 and intermittent engagements with 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 small discrete pieces of work spread out over weeks and months. Given these commitments, the day-to-day organisation of her work often required significant context switching across several different collaborative engagements. As P7 discusses, this context switching brings enormous challenges and requires considerable work to manage:

“The marketing piece of the business is so unique in the fact that a lot of our projects are very small. I have some projects that are only like 20 hours long and other projects that I'm on like that are Salesforce plus marketing plus these things. They can be thousands of hours, so because of the nature of the very small work, I'm on a lot of projects like I am looking. I'm not in charge of, but I am a part of about 15 to 20 different projects right now where some of my co-workers are on three. 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 just a lot to keep track of. So in our internal AI system Plaito, I can go in and pull up projects and every time I record a meeting I get like summary notes after and so I can talk to AI about my project and say hey, when was the last time I talked to this client about this subject and it'll say here were like all 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 all you would be like, yes, it was decided in this call. Here's like the chat transcript. Or no, it was determined that we talked about this later and that's how I can kind of remember, OK. I remember now when I talked to this one client about a very specific thing six months ago. I don't remember what we talked about, but now I'm able to like, talk to Plaito. They can keep me honest with what's already been talked about 'cause you don't want to go to a client and ask them the same things over and over and over again. They get annoyed. It's annoying. I would hate to be asked the same questions 1000 times. So having Plaito to be able to ask those questions to is great...I can ask Plaito about it and say hey for this specific client did they ever make a decision in our calls did they make a decision about this specific topic and then it'll pull up the transcriptions for me and...it'll say yeah they decided this and I'm like thank goodness because it's really big in consulting to not ask your clients the same questions a thousand times like - no one likes that. So, being able to get the answers on my own without having to rewatch recordings for hours on end, trying to make sure I'm finding the right recording, and being able to get the information that I need quickly is helpful" - P7 Business Analyst

What we see here is how in the sequential organisation of client interactions, P7's ability to invoke and demonstrate knowledge of prior encounters, discussions, and questions functions as a method for maintaining the recognisable orderliness of the professional relationships. Remembering, and also displaying that one remembers is not simply a cognitive accomplishment but also entails *relationship work* - accountable practices in which the participant orients to the relevancies of the ongoing client relationships – framing them as an continuous relationship rather than just a collection of disconnected calls and meetings. Significantly, such displays in their current conversations operate reflexively through explicit indexical reference to what was discussed in past interactions. As she highlights, forgetting or failing to acknowledge earlier discussions would risk being seen as inattentive, unprofessional, or even disrespectful. Through the act of remembering and weaving past discussions into present conversations, the participant was able to engage in a moral performance of care, competence and continuity with the client. It is through these practices that they were able to demonstrate their professionalism to a client – a key aspect of consulting work.

Participant P7 highlighted how this was previously challenging. Prior to the introduction of Plaito, preparation for these ongoing client encounters would have taken much longer potentially requiring hours searching for the right meetings, trawling through recordings and reviewing notes. Now, she recounted how she was able to question Plaito directly about the content of previous discussions with clients and whether certain things have been discussed and decided and when. 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^[7] Project Planning Meeting that took place on 2024-07-26...

P7: “Can you summarise any meeting where MCAE or Pardot^[8] 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...”

As we can see from the excerpt, through a quick exchange with Plaito, she can quickly obtain the information that she needs to effectively perform the professional relationship work in current client meetings.

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, cold calls, you know, or 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're like, OK, this sounds interesting. And then they book a meeting with me via like, my automatic calendar link. And then it's like, well, I have meeting. OK, who booked me, which is great. It's a great thing to have... I generally know who I'm booking the meeting with, but if they're just, you know, putting time on my calendar, 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 entails call preparation work to understand the client, the company, their role in the business, and what they care about. Such understanding allows them to approach the conversation in the right way. Prior to the introduction of Plaito, such preparation work would take considerable time with the need to build in structure to the daily schedule to try and accommodate it – but often leaving them feeling rushed.

“Yeah, I feel a little less rushed, you know, to be prepared for the meetings. You know, I usually would give myself like a 15-20 minute gap between any customer meeting just so that I could go and, like, find this person, find the company. Now I can like 5 minutes before the call. I can hit the “Call prep” button

and I can just do a quick review of everything about this person and then head into the call and confidence and be able to, you know, speak to some key pieces... What the call Prep does is looks at the person that is on my calendar, invite looks at their LinkedIn and it will give me just a high level of you know, where they currently work, some of their history, you know, some quick key links to the company website - their actual LinkedIn profile if I want to dive in a bit more. But just gives me an idea on like who I'm going to be speaking to, because sometimes, you know with sending out a lot of emails, people will book a time with you and you're like, OK, I need to know who booked this from which sequence, who are they, what do I need to be prepared, you know, when I jump on this call and you know, half an hour... and what matters to them, so those that's usually what I use call prep for is usually client-based communication... Obviously it's not going to be everything, but it's enough, you know, to have a warm conversation. So I can just focus on actually talking to the clients because then I can get that information so much faster that there's no hesitation of, you know, oh, I have to make sure I know all about this client before I reach out to them. So I do feel like 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 are able to 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 pre-call activities:

“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... Some customers I think they have the notion that you're the only people they work for.” P1 Principal Consultant

5 DISCUSSION

In this paper, we have presented a case study of evolving work practices of knowledge workers in an IT services organisation following the introduction of an organisationally tailored LLM agent, grounded in the systems, tools and documents comprising the organisational knowledge base and practices. Building on a growing corpus of insightful research into the impacts of Gen AI on knowledge workers, we extend these efforts to consider more the impact of an organisation-wide deployment of these tools on the collaborative organisation of knowledge work. Our concerns here were not so much with the idea of the agent as a collaborating partner (albeit while acknowledging its agency) but rather, with how the LLM

agent became actively implicated in an evolving set of collaborative practices of knowledge work within the organisation. The focus on a single organisation offers several important benefits to these concerns. In the first instance, the organisation as an entity offers a common set of established organisational goals and practices that bring a coherence to the working activities under scrutiny. Many of the participants work on common projects and client relationships while bringing the perspective of different types of stakeholders. Secondly, with the system adapted and grounded in organisational information and systems, it represents an important opportunity to understand its use beyond the more generic and organisationally isolated off-the shelf offerings. Thirdly, as an organisation wide deployment, there may be significant network benefits and influences on collective working practices.

What is apparent in the fieldwork is how the utility and influence of these systems cannot be fully understood in isolation of their embedding within the wider organisational ecosystem of enterprise tools and knowledge. The explicit integration of the agent within the Slack communication and collaboration environment is an important example here. Such integration enabled the ability to leverage pre-existing organisational structures of client and project management that already shape collaborative practice in significant ways. These structures within Slack have further reach into other parts of the other enterprise, customer, employee and project management systems ecosystem offering an additional avenue through which they can be relevant in context- and client-specific ways. It is this context which seamlessly grounds the responses of the LLM agent in project-specific and client-specific ways. This leads to a broader notion of *agential ubiety* [cf. 5], which makes salient the idea of the agent being situated in a particular place and time. The concept of *ubiety* features in Buxton's [5] discussion of ubiquitous computing, smart devices and ecologies of devices to argue that intelligence and smartness are not embedded in devices per se but that intelligence and utility emerge from the fit, placement and appropriateness of the device in the right place and time and through their interrelationship with the ecosystem of other tools and artefacts. Extending these arguments, the notion of *agential ubiety* highlights an important need to consider how these LLM systems can amplify their value and agency through their specific and appropriate embedding in existing organisational ecosystems.

Of further significance in the work was the material transformation of meetings through their recording, transcription and summarisation. There are echoes here of Erickson and colleagues' discussion of text-based computer-mediated conversation which functioned both as communication *and* persistent documentation [e.g. 15, 16, 17] that could form part of an organisational knowledge base. In the current context we saw how online video and hybrid meetings were made material and persistent. But in terms of understanding the impact of these tools on collaboration in knowledge work, their implications extend beyond the concept of passive documentation of existing collaborative practice. Rather we need to acknowledge a deeper agency of these systems in shaping collaboration in important ways as people actively adapt their behaviour in anticipation of the prospective capabilities of the *persistent meeting*. In the first instance, there were important ways in which people adapted their communication style to make themselves more legible to the AI in the meeting, highlighting opportunities where feedback might enhance the ability to do so. We saw these adaptations manifest in valuable ways through enhanced ability to attend to and engage with clients during meetings which can be such a significant component

of building trusted client relationships in service-oriented industries. Notably, this extends beyond a straightforward offloading of the cognitive demands of notetaking to an automated agent. Rather, as people become more familiar with these tools, and the prospective ability to interrogate, *remix* and repurpose the persistent meeting, they develop new opportunities to adapt in-the-moment conversation and collaboration to more closely align with the social demands of the encounter while maintaining the standardised documentation demands of their work – for example, where interview conversations no longer needed to rigidly follow a structured question list in the moment because such a structure could be introduced later through agent-based *remixing* of the meeting document.

Fully understanding the impact of these tools in knowledge work and how we design for them, requires us to give greater consideration to their role along trajectories of collaborative project work. These tools can be conceived as a dynamic and personalised form of “meeting bridging” [cf. 56] that link the synchronous work of meetings to other forms of synchronous and asynchronous collaboration and documentation. Furthermore, we need to think of the effects of these collaborative trajectories in bi-directional terms where there are both prospective and retrospective influences on the in-situ organisation of collaborative work. With the materialisation of meetings, the possibilities to search and interrogate them in conversation with the LLM agent allows the transcripts to be remixed into new collaboration and coordination artefacts. For example, templated documents of frequently performed projects activities are now actively brought into a meeting to structure the talk in ways that will facilitate the AI to more effectively fill in the blanks later. Meetings and conversation, in this regard, were structured in such a way as to *talk a document into existence* – asking questions in ways that would “*lead the [collaborative] witness*” and enable the recorded transcription to be appropriately parsed through subsequent prompting. Such behaviour was grounded in existing organisational document templates, again highlighting how the utility of these agents is bound up in their ability to effectively leverage and amplify existing and familiar organisational coordination artefacts. But what was also notable was how the participants created new document templates to exploit the new ability to talk documents into existence through their meetings. Through the appropriation of familiar enterprise productivity software and agent prompting, participants were empowered to develop their own improvised semi-automated workflows. Of importance here is the bottom-up nature of these behaviours. Rather than imposing workflows from above (albeit these sometimes being important), the participants were able to engage in more DIY approaches with familiar tools and in ways that aligned with the specific demands of their situated context.

Other enablers of effective collaboration across multiple projects trajectories related to the practices of re-establishing project context and maintaining situational awareness. Facilitating these practices were a new form of *managerial umwelt* enabled by the aggregative and temporal surfacing capabilities of the LLM agent to pull together information from disparate project related sources such as customer, project and employee management systems as well as across different temporal points in the project. The significance of this for the organisation can be seen in the ways that this aggregation behaviour was encoded in prompt shortcut buttons on the LLM agent, that offered frequently accessed standardised views for summarising key elements of project status. Access to these consolidations is quickly made a

routine part of daily practice, in particular for managers coordinating activity across multiple projects and teams. Empowering organisations to construct these around the repetitive features of their own work would be an important consideration to support. Over and above the creation of standardised organisational information aggregations, the establishment of situational awareness and context was further supported through more *interrogative aggregation* in which contextually tailored prompting was used to surface relevant information from across the siloed knowledge bases and enterprise systems to address the in-situ needs of collaborative knowledge work (such as, for example, supporting pre-meeting preparatory work that facilitates the presentation a professional competency to clients).

Of further interest to the concerns of situational awareness and coordination practices are some of the more innovative possibilities enabled through the LLM analysis of materialised meetings. The evolving practices that developed around sentiment analysis scores are an important exemplar here, highlighting new ways in which collaborative knowledge work activity can be analysed, understood and visualised as an input to the negotiated order of work. As has been noted in previous CSCW discussions, often relatively simple abstracted representations about work activity can be *socially translucent* and informative for collaboration when considered in the midst of richer contextual understanding [e.g. 15, 16, 36, 26, 27,32]. With this in mind, the capabilities of the LLM agent to extract and analyse communication and collaboration characteristics of workplace encounters points to potential new ways in which these events can be meaningfully understood and represented in socially translucent ways and around which new collaborative practices can be built. The findings highlighted how these agent-generated representations promoted visibility of certain types of activity that could initiate collaborative enquiry to seek underlying explanation for particular outcomes. They also demonstrated how making such characterisations visible within organisational structure can bring behavioural accountability. This in turn can have important influences on the ways that in-situ collaboration is organised with a view to achieving visibly positive sentiment. While these insights offer potential inspiration for other forms of socially translucent uses of agentic analysis, caution must also be taken to avoid the potential dehumanising pitfalls of algorithmic monitoring and management in the knowledge workplace. The tensions between workplace awareness, monitoring and surveillance is a longstanding one within CSCW since the very early days of media space research [e.g.21]. Dealing with this tension is as much about organisational culture as it is about system design. In the current study, such signals were seen as a valuable for enquiry and collective understanding of opportunities for intervention rather than a system to monitor and chastise.

The findings have also highlighted more fundamental ways in which workplace collaborations can be usefully reconfigured through creative interactions with the LLM agent to avoid coordination burden and make more cost-effective use of resource expertise. For example, we saw how the LLM agent supported extraction of information requirements from subject matter experts, that were then used to structure a client conversation in their absence, and in which linguistic “breadcrumbs” were deliberately left in the meeting, to be subsequently picked up post-meeting meeting from the recorded transcript through interrogation of the LLM agent. In this way, the need for synchronous presence of multiple and expensive subject matter experts in a long client meeting was avoided with a shift to more asynchronous

engagement – reducing the time commitment and coordination requirements of a synchronous and colocated meeting. This points to a broader set of opportunities and considerations for how we understand the significance of these systems in the organisation of collaborative knowledge work and the evolving ways that we come to configure synchronous and asynchronous elements together with colocated and remote. As organisations continue to grapple with these concerns in their everyday operations and structures, there remains a strong need for more ongoing research into these issues to help us design for a better future of work.

Declarations

The authors declare no competing interests.

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.

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

1. Asthana, S. Hilleli, S., He, P. and Halfaker, A. (2025) Summaries, Highlights, and Action Items: Design, Implementation and Evaluation of an LLM-powered Meeting Recap System. In *Proceedings ACM Human-Computer Interaction*, 9(2), Article CSCW176 (May 2025), 29 pages.
2. Barad, K. (2007) Meeting the Universe Halfway: Quantum Physics and the Entanglement of Matter and Meaning. London: Duke University Press.

3. Brachman, M., El-Ashry, A., Dugan, C. and Geyer, W. (2024) How Knowledge Workers Use and Want to Use LLMs in an Enterprise Context. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, Article 189, 1–8.
4. Bridle, J. (2023) *Ways of Being Animals, Plants, Machines: The Search for a Planetary Intelligence*. New York: Penguin Books.
5. Buxton, W. (2019) Ubiquity, Mobility & Transitions: The Quest for Ecological Intelligence. In *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces (ISS '19)*. Association for Computing Machinery, New York, NY, USA.
6. Bødker, S. (1989) A human activity approach to user interfaces. In *Human-Computer Interaction*, 4(3), 171–195.
7. Braun, V., & Clarke, V. (2006) Using thematic analysis in psychology. In *Qualitative Research in Psychology*, 3(2), 77–101.
8. Brynjolfsson et al (2023) Erik Brynjolfsson, Danielle Li, Lindsey Raymond (2025) Generative AI at Work. In *Quarterly Journal of Economics*, Volume 140(2), May 2025, Pages 889–942.
9. Butler, J., Jaffe, S., Baym, N., Czerwinski, M., Iqbal, S. Nowak, K., Rintel, S., Sellen, A., Vorvoreanu, M., Hecht, B., and Teevan, J. (2023) *Microsoft New Future of Work Report 2023*. Microsoft.
10. Butler, J., Vorvoreanu, M., Janssen, R., Sellen, A., Immorlica, N., Troy, A., Sarkar, A., Farach, A., Chouldechova, A., Olteanu, A., Cambon, A., Radhakrishna, A., Roseway, A., Zorn, B., Hecht, B., Goldstein, D., Joshi, D., Cutrell, E., and Kiciman, E., Ramos, G., Soares, G., Wallach, H. Drosos, I. Williams, J., O'Neill, J., Hofman, J., Teevan, J., Hernandez, J., Vaughan, J., Suh, J., Tang, J., Edwards, J., Bali, K., Inkpen, K., Madhavan, K., Bulman, L., Reicherts, L., Tankelevitch, L., Yang, L., Mott, M., Ochieng, M., Muchai, M., Baym, N., Abdulhamid, N., Marquardt, N., Hinckley, K., Bentley, M., Brown, D., Romat, H., Riche, N., Maina, S., Iqbal, S., Lindley, S., Nyairo, S., Blodgett, S., Gulwani, S. Sitaram, S. and Le, V. (2024) *Microsoft New Future of Work Report 2024*. Microsoft.
11. Chui, M., Hazan, E., Roberts, R., Singla, A., & Smaje, K. (2023). *The economic potential of generative AI: The Next Productivity Frontier*. McKinsey and Company
12. Dell'Acqua, Fabrizio, Edward McFowland III, Ethan Mollick, Hila Lifshitz-Assaf, Katherine C. Kellogg, Saran Rajendran, Lisa Kraye, François Candelon, and Karim R. Lakhani (2023) Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. In *Harvard Business School Working Paper, No. 24 – 013*, September 2023.
13. Duan, W., Flathmann, C., McNeese, N., Scalia, M., Zhang, R., Gorman, J., Freeman, G., Zhou, S., Hauptman, A. and Yin, X. (2025) Trusting Autonomous Teammates in Human-AI Teams - A Literature Review. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 1102, 1–23.
14. Duan, W., Zhou, S., Scalia, M., Yin, X., Weng, N., Zhang, R., Freeman, G., McNeese, N., Gorman, J., and Micael Tolston, M. (2024) Understanding the Evolvment of Trust Over Time within Human-AI

- Teams. In *Proceedings of ACM Human-Computer Interaction 8, CSCW2, Article 521* (November 2024), 31 pages.
15. Erickson, T., Halverson, C., Kellogg, W., Laff, M., and Vetting Wolf, T. (2001) Social Translucence Designing Social Infrastructures That Make Collective Activity Visible. In *Communications of the ACM*, 45. 10.1145/505248.505270.
 16. Erickson, T. & Smith, D., Kellogg, W., Laff, M., Richards, J. and Bradner, E. (1999). Socially Translucent Systems: Social Proxies, Persistent Conversation, and the Design of “Babble”. In *Proceedings of CHI '99, Conference on Human Factors in Computing Systems*, Pittsburgh, PA, USA.
 17. Erickson, T. and S. Herring, S. (2004) "Persistent conversation: a dialog between research and design,". In *Proceedings of 37th Annual Hawaii International Conference on System Sciences*, 2004. Big Island, HI, USA.
 18. Flathmann, C. Duan, W., McNeese, N., Hauptman, A., and Zhang, R. (2024) Empirically Understanding the Potential Impacts and Process of Social Influence in Human-AI Teams. In *Proceedings of ACM Human-Computer Interaction 8, CSCW1, Article 49* (April 2024), 32 pages.
 19. Guo, J., Mohanty, V., Piazzentin Ono, J., Hao, H., Gou, L. and Ren, L. (2024) Investigating Interaction Modes and User Agency in Human-LLM Collaboration for Domain-Specific Data Analysis. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, Article 203, 1–9.
 20. Gao, J. Guo, Y., Lim, G., Zhang, T., Zhang, Z. Jia-Jun Li, JT. and Perrault. S. (2024) CollabCoder: A Lower-barrier, Rigorous Workflow for Inductive Collaborative Qualitative Analysis with Large Language Models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 11, 1–29.
 21. Gaver, W., Moran, T. MacLean, A., Löfstrand, L., Dourish, P. Carter, K. and Buxton, W. (1992) Realizing a video environment: EuroPARC's RAVE system. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '92)*. Association for Computing Machinery, New York, NY, USA, 27–35.
 22. Hasker, S. (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>
 23. Hauptman, A., Duan, W. and McNeese, N. (2022) The Components of Trust for Collaborating with AI Colleagues. In *Companion Publication of the 2022 Conference on Computer Supported Cooperative Work and Social Computing (CSCW'22 Companion)*. Association for Computing Machinery, New York, NY, USA, 72–75.
 24. He, J. Houde, S., Gonzalez, G. E., Andrés Silva Moran, D., Ross, S.I., Muller, M. and D. Weisz, J.D. (2024) AI and the Future of Collaborative Work: Group Ideation with an LLM in a Virtual Canvas. In *Proceedings of the 3rd Annual Meeting of the Symposium on Human-Computer Interaction for Work (CHIWORK '24)*. Association for Computing Machinery, New York, NY, USA, Article 9, 1–14.

25. Hughes, J. A., Randall, D., & Shapiro, D. (1992, December). Faltering from ethnography to design. In *Proceedings of CSCW 1992 ACM conference on Computer-Supported Cooperative Work* (pp. 115–122).
26. Kaye, J., Levitt, M., Nevins, J., Golden, J., and Schmidt, V. (2005) Communicating intimacy one bit at a time. In *Extended Abstracts of CHI '05 conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1529–1532.
27. Kaye, J. (2006) I just clicked to say I love you: rich evaluations of minimal communication. In *CHI '06 Extended Abstracts on Human Factors in Computing Systems (CHI EA '06)*. Association for Computing Machinery, New York, NY, USA, 363–368.
28. Lauro, Q., Bigham, J., and Kotturi, Y. (2024) Exploring the Role of Social Support When Integrating Generative AI in Small Business Workflows. In *Companion Publication of the 2024 Conference on Computer-Supported Cooperative Work and Social Computing (CSCW Companion '24)*. Association for Computing Machinery, New York, NY, USA, 485–492. <https://doi.org/10.1145/3678884.3681895>
29. Lee, H-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, R., Banks, R., and Wilson, N. (2025) The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects from a Survey of Knowledge Workers. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 1121, 1–22.
30. Lindley, S. and Wilkins, D.J. (2023) Building Knowledge through Action: Considerations for Machine Learning in the Workplace. In *ACM Trans. Computer-Human Interaction*. 30(5), Article 72 (October 2023), 51 pages.
31. Muller, M., He, J. and Weisz, J. (2024) Workplace Everyday-Creativity through a Highly-Conversational UI to Large Language Models. *J. ACM* 37, 4, Article 111.
32. Nardi, B., Whittaker, S. and Bradner, E. (2000) Interaction and Outeraction: instant messaging in action. In *Proceedings of the 2000 ACM conference on Computer supported cooperative work (CSCW '00)*. Association for Computing Machinery, New York, NY, USA, 79–88. <https://doi.org/10.1145/358916.358975>
33. Nonaka, I. and Takeuchi, H. (1995) *The Knowledge Creating Company: How Japanese Companies Create Dynamics of Innovation*. Oxford: Oxford University Press.
34. Nonaka, I. (1994) A dynamic theory of organisational knowledge creation. In *Organisation Science* 5(1), 14–37.
35. Nonaka, I. and Toyama, R. (2015) The Knowledge Creating Theory Revisited: Knowledge Creation as a Syntheizing Process. In Edwards, J.S. (2015) *The Essentials of Knowledge Management*. Basingstoke: Palgrave Macmillan.
36. O'Hara, K. Massimi, M., Harper, R., Rubens, and Morris, J. (2014) Everyday dwelling with WhatsApp. In *Proceedings of the 17th ACM conference on Computer supported cooperative work and social computing (CSCW '14)*. Association for Computing Machinery, New York, NY, USA, 1131–1143. <https://doi.org/10.1145/2531602.2531679>

37. Orlikowski, W.J. (2002) Knowing in Practice: Enacting a Collective Capability in Distributed Organizing. In *Organization Science* 13(3) (May 2002), 249–273.
38. Orlikowski, W.J. (2006). Material knowing: the scaffolding of human knowledgeability. In *European Journal of Information Systems*, 15, 460–466.
39. Russo. D. (2024) Navigating the Complexity of Generative AI Adoption in Software Engineering. In *ACM Trans. Softw. Eng. Methodol.* 33(5).
40. Sako. M. (2024) How Generative AI Fits into Knowledge Work. In *Communications of the ACM* 67(4) (April 2024), 20–22. <https://doi.org/10.1145/3638567>
41. Schmidt, K., Bannon, L. (1992) Taking CSCW seriously. In *Computer-Supported Cooperative Work* 1, 7–40 (1992).
42. Schmutz, J.B. Outland, N., Kerstan, S., Georganta, E. and Anna-Sophie Ulfert, A-S (2024) AI-teaming: Redefining collaboration in the digital era. In *Current Opinion in Psychology*, 58.
43. Seely Brown, John & Duguid, Paul. (2000). *The Social Life of Information*. Boston: Harvard Business Review Press.
44. Sellen, A and Harper. R. (2003) *The Myth of the Paperless Office*. MIT Press, Cambridge, MA, USA.
45. Shin, J., Koch, J., Lucero, A., Dalsgaard, P. and Wendy E. Mackay, W. (2023) Integrating AI in Human-Human Collaborative Ideation. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23)*. Association for Computing Machinery, New York, NY, USA, Article 355, 1–5.
46. Simkute, A., Tankelevitch, L. & Kewenig, V. & Scott, A., Sellen, A. & Rintel, S. (2024). Ironies of Generative AI: Understanding and Mitigating Productivity Loss in Human-AI Interaction. In *International Journal of Human-Computer Interaction*. 1–22. 10.1080/10447318.2024.2405782.
47. Star, S.L., and Strauss, A (1999): Layers of Silence, Arenas of Voice: The Ecology of Visible and Invisible Work. In *Computer Supported Cooperative Work (CSCW)*, 8(1–2). DOI: 10.1023/A:1008651105359. Springer. PISSN: 1573–7551. pp. 9–30.
48. Strauss, A. (1988). The Articulation of Project Work: An Organizational Process. In *The Sociological Quarterly*, 29(2), 163–178.
49. Strauss, A. (1985) *The Social Organization of Medical Work*. Chicago: University of Chicago Press.
50. Suchman, L. (1994) Supporting Articulation Work. In *Proceedings of the 5th International Conference on Women, Work and Computerization: Breaking Old Boundaries, Building New Forms*, Manchester, UK.
51. Sun, Y., Jang, E., Ma, F. and Wang. T. (2024) Generative AI in the Wild: Prospects, Challenges, and Strategies. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 747, 1–16. <https://doi.org/10.1145/3613904.3642160>
52. Vanukuru, R., Panda, P., Chen, X., Scott, A. Tankelevitch, L. Rintel, S. (2025) Designing Interfaces that Support Temporal Work Across Meetings with Generative AI. In *Proceedings of DIS 25*, Funchal,

Portugal.

53. Verheijden, M. and Mathias Funk. M. (2023) Collaborative Diffusion: Boosting Designerly Co-Creation with Generative AI. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23)*. Association for Computing Machinery, New York, NY, USA, Article 73, 1–8.
54. von Uexküll, J. (2010) *A Foray into the Worlds of Animals and Humans: With a Theory of Meaning*. Univ Of Minnesota Press.
55. Wagman, K., Dearing, M. and Chetty, M. (2025) Generative AI Uses and Risks for Knowledge Workers in a Science Organization. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 1199, 1–17.
56. Wang, R., Qiu, L., Cranshaw, J., and Zhang. A. (2024) Meeting Bridges: Designing Information Artifacts that Bridge from Synchronous Meetings to Asynchronous Collaboration. In *Proceedings ACM Human-Computer Interaction. 8, CSCW1, Article 35* (April 2024), 29 pages.
<https://doi.org/10.1145/3637312>
57. Wolf, C. & Blomberg, Jeanette, J. (2019). Evaluating the Promise of Human-Algorithm Collaborations in Everyday Work Practices. In *Proceedings of the ACM on Human-Computer Interaction. 3*. 1–23.
[10.1145/3359245](https://doi.org/10.1145/3359245).
58. Woodruff, A., Shelby, R., Gage Kelley, P. Rousso-Schindler, S., Smith-Loud, J. and Wilcox, L. (2024) How Knowledge Workers Think Generative AI Will (Not) Transform Their Industries. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 641, 1–26.
59. Yun, B., Feng, D., Chen, A., Nikzad, A. and Salehi, N. (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)*. Association for Computing Machinery, New York, NY, USA, Article 634, 1–19. <https://doi.org/10.1145/3706598.3713337>
60. Zhang, R., Duan, W., Flathmann, C., McNeese, N., Freeman, G., and Williams, A. (2023) Investigating AI Teammate Communication Strategies and Their Impact in Human-AI Teams for Effective Teamwork. In *Proceedings ACM Human-Computer Interaction 7, CSCW*, Article 281 (October 2023), 31 pages.

Footnotes

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. 16, 17]
2. We draw on Barad's [2] 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.

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.
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.
5. *Umwelt* is a term coined by Jakob von Uexküll [54]. The term refers to the subjective perceptual world of an individual organism – how that individual 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 in *Ways of Being* [4].
6. 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.
7. Marketing Cloud Account Engagement (MCAE) is a B2B marketing automation platform by Salesforce.
8. MCAE was formerly known as Pardot

Figures

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

[redacted] | Item Mgmt. Walkthrough

2024-10-22 11:57 AM Eastern

Attendees:

[redacted]
[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.

Figure 1

An example of a meeting summary generated by Plaito

File Edit View Insert Format Data Tools Extensions Help Accessibility					
100% 123 Default... 10 B I A					
A1 Epic					
	A	B	C	D	E
1	Epic	Story Name	User Story	Acceptance Criteria	Solution Components
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. - Users must be able to view and edit these fields on the Campaign detail page. - 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 2. Update the campaign page layout to include the new fields.
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. - Users should be able to easily add or remove child campaigns from a parent campaign. - 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.
4	Campaign Management	Connect to Opportunity so that the dollars can be tracked.	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. - When an Opportunity is created from a Campaign, the Campaign ID should be automatically populated in the Opportunity record. - The Opportunity record should display the Campaign Name for easy reference. - The Campaign's "Expected Revenue" field should be updated to reflect the sum of the amounts from all related Opportunities. - 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
CRM User Stories CPQ User Stories Service Cloud User Stories Integration Stories Data Stories					

Figure 2

User Stories Spreadsheet in preparation for conversion to Jira tickets