

What Did You Do Now?

A Qualitative Case Study on Knowledge Worker' Interactions with Agentic AI

Oscar Strömberg (24887) and Joel Edholm (24776)

Abstract: This study draws on 22 semi-structured interviews with knowledge workers using an AAI-powered web browser across diverse professional contexts (e.g., consulting, engineering, marketing, and operations) to explore how AAI reshapes goal-oriented action in knowledge work. Employing Affordance Theory as our theoretical lens, we show how workers navigate a complex landscape in which the same technological features can both enable and constrain goal attainment, depending on the situated alignment between workers' goals and technological capabilities. We identify four affordances (e.g., automating workflows) and four constraints (e.g., increasing idle time) that emerge from knowledge workers' interactions with the material technological features of AAI. Further, our findings indicate that the likelihood of actualizing action potentials is shaped by contextual factors at the individual (e.g., AAI literacy), task (e.g., stakes and scale), and organizational (e.g., organizational ways of working) levels. This study offers one of the first empirical investigations of autonomous AI agents in knowledge work and contributes an empirically derived conceptual model of knowledge worker–AAI interactions. This contributes to a broader understanding of how knowledge workers' interactions with autonomous technologies generate more fundamentally agentic, rather than purely informational, action potentials, and illustrates how knowledge workers not only interrogate what AAI generates but also actively assess whether to delegate agency for action responsibilities – creating new accountability dynamics in knowledge work.

Keywords: Agentic Artificial Intelligence, Knowledge Work, Affordance Theory, Human-AI Interaction, Autonomous Action-Taking

Supervisor:	Elmira van den Broek
Date examined:	16/12/2025
Discussant(s):	Jasmin Hepp & Daniel Haab
Examiners:	Roberto Verganti & Constanze Eib

Acknowledgements

We want to thank Elmira van den Broek for her invaluable support and guidance throughout this thesis semester. Her profound knowledge of our chosen phenomenon proved immensely helpful during the writing period. We also want to express our gratitude to the peer group for reviewing our work and offering thoughtful suggestions.

Sincerely,

Oscar Strömberg & Joel Edholm

Table of Contents

1. Introduction	1
1.1 Background	1
1.2 Research Gap	2
1.3 Study Focus	3
1.4 Purpose & Expected Contribution	3
2. Theory	4
2.1 Literature Review	4
2.1.1 AI in Knowledge Work	4
2.1.2 Agentic AI: A New Frontier	7
2.2 Synthesis and Research Gap	9
2.3 Theoretical Lens: Affordance Theory	9
3. Methodology	12
3.1 Research Design and Approach	12
3.1.1 Scientific Research Approach	12
3.1.2 Research Case	13
3.1.3 Research Process	14
3.2 Data Collection	15
3.2.1 Phase I	15
3.2.2 Phase II	16
3.2.3 Interview process	17
3.3 Data Analysis	18
3.4 Quality of Research and Ethical Considerations	20
3.4.1 Discussion on Quality	20
4. Findings	21
4.1 Case Context	21
4.1.1 The Case Company: AgentCo	21
4.1.2 User Goals	22
4.1.3 AAI Features	23
4.2 AAI Affordances	25
4.2.1 Affordance 1: Automating Workflows	25
4.2.2 Affordance 2: Extracting Real-Time Web-Based Information	27
4.2.3 Affordance 3: Maintaining Contextual Continuity	28
4.2.4 Affordance 4: Distributing Specialized Labor	29
4.3 AAI Constraints	31
4.3.1 Constraint 1: Generating Inconsistent Memories	31
4.3.2 Constraint 2: Increasing Idle Time	33
4.3.3 Constraint 3: Struggling With Human-Trivial Tasks	34
4.3.4 Constraint 4: Overdelivering on Simple Tasks	35
4.4 Contextual Determinants of Affordance-Actualization	36

4.4.1 Individual Characteristics: Ways of Knowing and Need for Control	36
4.4.2 Task Characteristics: Stakes, Scale, and Error Tolerance	40
4.4.3 Organizational Context: Restrictive Policies and Ways of Working	42
5. Analysis	44
5.1 A Conceptual Model of Knowledge Worker–AAI Interactions	44
5.1.1 Affordances and Constraints	44
5.1.2 Affordance-Actualization Potential	46
6. Discussion	48
6.1 Answer to Research Question	48
6.2 Theoretical Contributions	49
6.3 Practical Implications	51
6.4 Limitations and Future Research	51
References	53
Appendices	62
Appendix A. Comparison of traditional AI and agentic AI capabilities	62
Appendix B. Interview Guide: Round 1	63
Appendix C. Interview Guide: Round 2	64
Appendix D. Summary of Affordances and Constraints: Grouped by AAI Feature	65
Appendix E. How AgentCo Works: Answered by an AgentCo Agent	66
Appendix F. Disclosure of AI Use	67

List of Figures

Figure 1. Visual Overview of Research Process	14
Figure 2. Coding and Data Analysis Excerpt	19
Figure 3. Google Trends – Search Terms: “AI browser“ and “AI agents“	22
Figure 4. AAI–affordances, constraints, and actualization in knowledge work	44

List of Tables

Table 1. Table 1. Interview Sample: Phase I	15
Table 2. Interview Sample: Phase II	16
Table 3. Overview of Affordances	25
Table 4: Overview of Constraints	31

Definitions

Term	Definition
Knowledge workers	“The man who works exclusively or primarily with his hands is the one who is increasingly unproductive. Productive work in today’s society and economy is work that applies vision, knowledge and concepts – work that is based on the mind rather than on the hand.” - (Drucker, 1957)
Epistemic technologies	"[...] tools that play a central role in the ongoing construction of knowledge." - (Anthony, 2018)
Artificial Intelligence (AI)	"It is the science and engineering of making intelligent machines, especially intelligent computer programs." - (McCarthy, 2007)
Generative Artificial Intelligence (GenAI)	"The term generative AI refers to computational techniques that are capable of generating seemingly new, meaningful content such as text, images, or audio from training data." - (Feuerriegel et al., 2024)
Agentic Artificial Intelligence (AAI)	"Agentic AI refers to autonomous, goal-driven systems that can operate on their own for long periods, requiring minimal human supervision. Unlike traditional AI or standard large language models (LLMs) that respond only to single prompts, these systems can understand broad objectives, break them down into smaller tasks, and carry out multi-step plans while adapting to feedback from changing environments. Their "agentic" nature also highlights their ability to take on responsibilities from humans, act purposefully, and be accountable for the results they produce." - (Bandi et al., 2025)
Affordance	“The concept of technology affordance refers to an action potential, that is, to what an individual or organization with a particular purpose can do with a technology or information system.” - (Majchrzak & Markus (2012)
Constraint	“Technology constraint refers to ways in which an individual or organization can be held back from accomplishing a particular goal when using a technology or system.” - (Majchrzak & Markus (2012)

1. Introduction

1.1 Background

“AI agents will become the primary way we interact with computers in the future.”

Satya Nadella, CEO, Microsoft

The emergence of Agentic Artificial Intelligence (AAI) marks a significant evolution at the intersection between algorithmic technologies and knowledge workers – those who *“think for a living”* (Drucker, 1957). Defined as *“AI systems that can pursue complex goals with limited direct supervision”* (Shavit et al., 2023, p.1), emerging AAI technologies have the potential to not only reason but also take action on behalf of their users (Acharya et al., 2025). A recent report, surveying 10,000 knowledge workers, shows that they spend 60% of their time on *“work about work”* – such as looking for information and switching between tools – rather than doing meaningful work (Talbert, 2025). Hence, even in the age of Artificial Intelligence (AI), where algorithms can generate predictions and artefacts that emulate those of knowledge workers (Faraj et al., 2018), they remain unable to take action. One could say we have *“created a world of brilliant advisors who can’t lift a finger to help”* (Bornet et al., 2025, p.2).

McKinsey & Company recently posited that AAI could play an active role in solving this issue by shifting technologies from passive tools to active companions (Sukharevsky et al., 2025). As such, the diffusion of AAI across professional settings represents a potentially pivotal development for knowledge workers, enabling them to outsource execution responsibilities to autonomous agents (Bornet et al., 2025). Optimism about AAI’s potential is reflected in its rapid diffusion in professional settings, with Gartner predicting that 33% of enterprise software applications will incorporate AAI by 2028 (Coshov, 2024). While it is tempting to envision a utopian future in which AAI operates as digital co-workers capable of executing tasks autonomously, studies on knowledge workers’ use and interactions with previous generations of AI technologies suggest the reality might be more nuanced.

By conceptualizing AI as an *“epistemic technology”* (Anthony, 2018, p. 661) involved in the creation of knowledge, scholars have inferred that AI and knowledge workers are intimately related (e.g., van den Broek, 2025) – since knowledge workers are individuals primarily

concerned with creating and using knowledge (Davenport, 2005). The intersection of AI and knowledge work has therefore attracted particular scholarly interest in recent years (e.g., Faraj et al., 2018; Anthony et al., 2023; van den Broek, 2025). As these technologies enter professional settings, Pakarinen & Huising (2023) argue that they become entangled in the situated relational practices through which knowledge workers generate knowledge. With this development, challenges related to differences between algorithmic and human “*ways of knowing*” have become evident (Faraj et al., 2018, p.65). For example, the reasoning behind AI’s knowledge claims is often characterized as a “*black box*” (Christin 2020; Anthony et al., 2023), where humans are unable to understand how the algorithm reached a particular conclusion. When encountering this opacity of algorithmic reasoning, knowledge workers have been found to engage in practices of interrogating (Lebovitz et al., 2022), questioning (Anthony, 2018), and translating (Waardenburg et al., 2022) algorithmic output to make sense of AI’s knowledge claims. These practices reveal how the material properties of AI (e.g., the opacity of AI outputs) shape the actions knowledge workers take – showing the “*material agency*” and emphasizing the role of technologies themselves in knowledge work (Newell, 2015, p.9). Scholars have therefore called for the need to increasingly consider the material properties of AI technologies themselves in studies of IS use and knowledge work (Newell, 2015; Baird & Maruping, 2021)

1.2 Research Gap

The introduction of AAI represents a fundamental shift towards increasingly sophisticated autonomous agents (Krakowski, 2025), encompassing unique technological properties that enable them to take actions on behalf of users across increasingly complex tasks (Bornet et al., 2025). In a recent MISQ editorial paper, Leonardi (2023) articulates that, as AI increasingly acts and adapts without human intervention, changes in work will depend on the configuration of agency between humans and AI, in ways that enable and constrain action – presenting novel implications for “*who or what can act, and how*” (p.17). As individuals practice agency in their construction of knowledge (Prawat, 1996), and knowledge workers’ primary concern lies in constructing, transmitting, and using knowledge (Davenport, 2005), the recursive redistribution of agency between humans and AAI (Leonardi, 2025) is likely to impact how these workers interact with these novel and autonomous technologies. While prior research on AI in knowledge work offers valuable insights into the interactions between skilled workers and epistemic technologies, these theoretical concepts are based on earlier

generations of AI technologies that primarily generate recommendations and artefacts for knowledge workers to evaluate and act upon – and may therefore be limited in their applicability to AAI technologies that can autonomously execute actions on workers’ behalf. Despite initial conceptual interest in AAI, empirical research examining how knowledge workers interact with these systems in their real-life everyday work remains, to the best of our knowledge, scarce.

1.3 Study Focus

Accordingly, this study aims to address the unexplored and timely gap in understanding how the diffusion of AAI technologies creates new configurations of human-AI interactions in knowledge work. The following research question guides the study:

How do knowledge workers interact with Agentic AI to achieve work-related goals?

To answer this question, we employ a qualitative single-case research design, studying AgentCo: a startup that has built an AAI-powered web browser to empower knowledge workers through its agentic browsing capabilities. Through 22 semi-structured interviews with knowledge workers who use AgentCo’s AAI tool across diverse professional contexts – including, e.g., consulting, marketing, sales, and engineering – we explore how these professionals interact with AAI to achieve their work-related goals. Using Affordance Theory (Gibson, 1979) as a theoretical lens, we examine how knowledge workers’ interactions with the unique technological features of AgentCo’s AAI tool enable or constrain potential actions aimed at achieving particular work-related goals (Majchrzak & Markus, 2012).

1.4 Purpose & Expected Contribution

This study aims to make two primary contributions. First, we strive to provide empirical insights into an emerging and under-researched phenomenon, offering one of the first empirical investigations of AAI use in knowledge work contexts (Leonardi, 2023). Second, recognizing that the action possibilities of technology use are situated and relational (Leonardi, 2011; Majchrzak & Markus, 2012), we aim to contribute to the understanding of how situated relationships between AAI technologies and humans enable or constrain goal-oriented actions in knowledge work. In doing so, we aim to offer guidance to technology designers and users of AAI on how to facilitate better use of these novel technologies (Volkoff & Strong, 2017).

2. Theory

This chapter reviews prior literature on knowledge workers' interactions with earlier generations of AI technologies, then examines how Agentic AI differs through its novel technological features (2.1). Drawing on this literature, we identify a key research gap (2.2) before introducing the theoretical lens that guides our study (2.3).

2.1 Literature Review

2.1.1 AI in Knowledge Work

Advances in Artificial Intelligence (AI) have enabled algorithmic technologies to perform cognitive tasks once exclusive to skilled professionals (Loscher & Bader, 2022). Faraj et al. (2018) define these systems as emerging technologies that use machine learning to produce analyses and predictions that “*resemble those of a knowledge worker*” (p.1) – implying that the line between machine and human capabilities is blurring (Berente et al., 2021). Knowledge workers are individuals who think for a living (Drucker, 1957) and whose primary professional purpose revolves around the “*creation, distribution, or application of knowledge*” (Davenport, 2005, p.9). Contemporary AI has been closely linked to knowledge work through its role as an epistemic technology, or “*tools that play a key part in the ongoing construction of knowledge*” (Anthony, 2018, p. 661). While knowledge workers have long used algorithmic tools in their work, recent AI advances have significantly expanded the range of tasks these systems can perform (Allen & Choudhury, 2021). Understanding how knowledge workers interact with systems capable of resembling their own work has therefore become a critical focus for scholars examining the future of expertise, professional authority, and organizing (e.g., Kim et al., 2025; Pakarinen & Huising, 2023; Dell’Acqua et al., 2023; van den Broek, 2025; Anthony, 2018).

AI technologies, unlike previous technological advancements, are argued to present new threats to knowledge workers’ authority by emulating and challenging their expertise (Pakarinen & Huising, 2023). It is not uncommon to portray AI as systems capable of producing knowledge autonomously and to claim they can exceed human expertise by learning from datasets (Agrawal et al., 2022; Domingos, 2015). For the first time, roles that depend on cognitive skills are seen as vulnerable to being replaced by AI (Brynjolfsson & McAfee, 2014). However, recent scholars have offered alternative perspectives to these

claims. Pakarinen and Huising (2023) argue that those who view AI as a direct threat to workers' authority employ an overly substantialist view of expertise as an "*intellectual possession*" (p.2054) separate from its surrounding context. Exploring the situated relational practices through which knowledge is created, these scholars suggest that AI will become integrated into the network of interactions that shape knowledge rather than replace it (Pakarinen & Huising, 2023). Pettersen (2019) focuses, instead, on the problems in which professional knowledge is applied, positing that "*a great deal of knowledge work concerns highly complex problem solving and must be understood in contextual, social and relational terms*" (p.1) – suggesting knowledge work constitutes less generic characteristics, making it difficult for AI to replace. Faraj et al. (2018) offer another pragmatic view of the issue, suggesting that when human knowledge claims prove inferior to those of algorithms, these workers' attention will likely be redirected towards supervising, using, and improving algorithms.

Notwithstanding the ongoing debate on AI's potential to erode knowledge workers' expertise, these technologies are inevitably integrated across a broad range of knowledge-intensive contexts. For example, human resources (HR) professionals use analytical AI to assist with CV screening (Kaplan & Haenlein, 2019); loan consultants use predictive AI to establish loan terms (Strich et al., 2021); and radiologists take help from recognition algorithms to detect heart and lung diseases (Mukherjee, 2017). These instances all relate to decision-making that was previously grounded solely in knowledge workers' expertise, suggesting that AI should not be understood merely as a tool but as an active participant in redefining the nature of knowledge workers' expertise and introducing new challenges for professions to confront (Faraj et al., 2018).

Multiple studies highlight challenges of integrating AI technologies in knowledge work, frequently attributing these to differences in the making and evaluation of claims between algorithms and experts (Van den Broek et al., 2021; Pachidi et al., 2020). A recent paper by Kim et al. (2025, p.1) contrasts "*situated*" and "*algorithmic*" modes of knowledge work, positing that experts leverage intuition and embodied skill in relation to a specific contextual challenge to produce knowledge. In contrast, algorithms follow predefined processes to do so (Kim et al., 2025). Since AI algorithms are unable to tailor their suggestions to tacit knowledge accessible only through interactions with stakeholders in situated contexts, they often struggle to produce outputs that workers understand and feel motivated to use

(Pakarinen & Huising, 2023). This issue has directed scholarly attention to the emergence of new routines and roles as knowledge workers increasingly encounter algorithmic knowledge claims at work (Anthony, 2018; Waardenburg et al., 2022; Lebovitz et al., 2022). For example, Kellogg et al. (2019) introduced the term “*algorithmic brokerage*,” which was later explored in a study on the implementation of learning algorithms in policing, where certain employees moved from being messengers to interpreters and, finally, to curators of algorithmic outputs (Waardenburg et al., 2022). However, the challenges of implementing AI in knowledge work extend beyond differences in the making and evaluation of knowledge claims. Dell’Acqua et al. (2023) explored the use of AI technologies by consultants, finding that the models’ performance varies dramatically and unpredictably across seemingly comparable tasks, creating what the authors refer to as the “*jagged Frontier*” of AI. Therefore, the benefits of AI technologies, albeit significant, are not immediately evident to their users, suggesting it may be challenging for knowledge workers to discern appropriate contexts for AI deployment (Dell’Acqua, 2023).

To extend and maintain expertise, knowledge workers have long been assumed to engage in questioning of technological output (Knorr Cetina, 1999). However, unlike traditional technologies, whose logic can be unpacked, the reasoning process behind AI algorithms is opaque, yielding algorithmic outputs that are inscrutable to their users (Faraj et al., 2018; Anthony, 2021). Contemporary AI technologies are therefore often characterized as “*black box*” technologies (Christin 2020; Anthony et al., 2023). The issues with black box technologies become pronounced when knowledge workers need to understand and defend knowledge claims made by algorithmic technologies (Anthony, 2021). Lebovitz et al. (2022) examined how these workers engage in interrogation practices – by “*relat[ing] their own knowledge claims to AI’s knowledge claims*” (p. 139) – to address the opacity of output produced by AI technologies. Pakarinen and Huising (2023) further argue that, since knowledge workers remain accountable for the use of algorithmic outcomes, they must also be able to explain and verify their claims – a challenge exaggerated by the “*lack of explainability in AI technologies*” (p.2069). To this point, Lebovitz et al. (2022) found that radiologists bearing diagnostic accountability are inclined to dismiss algorithmic knowledge claims when they cannot interrogate the reasoning behind them.

In understanding the consequences of black-boxed technologies, Anthony (2021) investigated how bankers assembled and interpreted data using algorithmic technologies. The author

found that, in groups where junior bankers assembled analyses and seniors interpreted them, both junior and senior bankers were inclined to rely on algorithmic outputs without understanding their underlying logic, leading to inconsistent knowledge outputs and constrained junior bankers' knowledge development (Anthony, 2021).

Knowledge can be understood as socially constructed through relational practices; however, AI technologies are increasingly embedded within these practices as active participants (Pakarinen & Huising, 2023). This raises questions about the role of these technologies in shaping knowledge work. Newell (2015) addresses this through the concept of “*material agency*” (p.9) – the capacity of technologies to influence human behavior and work practices. By acknowledging that AI possesses material agency, we, in line with Newell (2015), argue for recognizing the role of material properties themselves in knowledge work. While prior literature offers valuable insights into the effects of the material properties of earlier generations' technologies (e.g., Faraj et al., 2018; Waardenburg et al., 2022), these may be less applicable to newer generations. The most recent development in AI, Agentic Artificial Intelligence (AAI), constitutes a qualitative leap from previous generations of AI (Acharya et al., 2025) towards advanced autonomous agents (Krakowski, 2025). Empirical evidence on the implications of AAI in knowledge work remains nascent. The following section examines AAI's distinctive technological properties and how they differ from previous AI generations.

2.1.2 Agentic AI: A New Frontier

The development of AI has undergone two fundamental shifts in recent years. First, the emergence of Generative AI (GenAI) systems such as ChatGPT, capable of generating content from natural language (Brynjolfsson et al., 2023). Unlike predictive AI, which identifies patterns, GenAI enables the generation of text, images, and code, altering how knowledge workers interact with algorithms (Krakowski, 2025; Dell'Acqua, 2023). Second, the emergence of Agentic Artificial Intelligence (AAI), defined as “*AI systems that can pursue complex goals with limited direct supervision*” (Shavit et al., 2023, p.1). This definition captures the fundamental departure from earlier technologies, as it “*doesn't just generate insights – it takes action*” (Bornet et al., 2025, p.16). The combination of reasoning power from Large Language Models (LLMs) with unique technological architectures enables AAI to take performative action in increasingly complex environments¹ (Bornet et al., 2025;

¹ See Appendix A for a detailed comparison of traditional AI and AAI capabilities

Bousetouane, 2025). Synthesising the nascent literature on AAI, we identify four material properties or features that distinguish AAI from previous algorithmic technologies.

First, autonomous action-taking that enables real-world execution through tool use and API calls (Bornet et al., 2025; Bousetouane, 2025). For example, when giving an agent the prompt 'send a message to my team on Slack. Open the Slack app and then click new message,' an LLM can digest and understand the required sequence as it has been trained on countless similar instances (Bornet et al., 2025). This understanding is then translated into actual use of tools, such as calling the Slack API to send the message. Second, AAI is capable of autonomous reasoning that enables decision-making under uncertainty by breaking complex tasks into sequential sub-tasks and adapting without human intervention (Acharya et al., 2025; Bousetouane, 2025). Third, multi-agent coordination facilitates collaboration between specialized agents through agentic orchestration layers, enabling the completion of increasingly complex tasks (Acharya et al., 2025). Bornet et al. (2025) invite us to think about this feature as an orchestra, where *"each musician plays a different instrument, following their own sheet music, yet together, they create a harmonious performance"* (p.91). Instead of having a single system handle everything, we can direct subtasks to specialised agents. Lastly, persistent memory allows AAI systems to store historical interactions, user preferences, and domain-specific knowledge across contexts (Bornet et al., 2025). Contrary to current GenAI systems that require users to re-enter context continuously, AAI builds on previous interactions and retains contextual understanding over time (Bornet et al., 2025).

While agentic systems capable of independently interacting in dynamic environments have existed for years, they have traditionally required rule-based programming or precise training of machine-learning models, limiting their practical applicability for more complex problem-solving (Yee et al., 2024). Lyytinen et al. (2021) argue that the autonomy of earlier AI technologies was controlled by humans, where tasks were carried out according to how humans programmed the systems (Berente et al., 2021). However, scholars are increasingly recognizing the agentic nature of IS technologies and their ability to perform tasks without human direction (Baird & Maruping, 2021; Leonardi, 2023). A recent paper by Leonardi (2025, p.2) conceptualizes these dynamics as *"agency loops,"* in which the *"responsibility for action circulates among human and non-human actors,"* rather than being simply delegated to technologies. Leonardi (2023) argues that as AI becomes increasingly capable of acting and adapting without human intervention, changes in work will depend on how agency is configured between humans and AI within their organizational context, in ways that enable

and constrain action. Berente et al. (2021) similarly state, “*the intersection between humans and autonomous AI is perhaps the key managerial issue of our time*” (p.1440).

2.2 Synthesis and Research Gap

Prior research on AI in knowledge work demonstrates how professionals interact with material properties of previous generations of AI technologies (e.g., inscrutability; Faraj et al., 2018) by engaging in translation (Waardenburg et al., 2022), interrogation (Lebovitz et al., 2022), and interpretation (Anthony, 2021) of algorithmic knowledge claims. These interactions are all part of the situated practices through which knowledge is generated alongside technologies (Pakarinen & Huising, 2023). Although these studies offer valuable empirical and theoretical insights into how knowledge workers interact with AI technologies, they predominantly focus on earlier generations of AI that generate predictions and other outputs for users to engage with. The introduction of AAI represents a fundamental shift towards increasingly autonomous agents (Krakowski, 2025), capable of acting independently without human intervention (Bornet et al., 2025). Despite initial conceptual interest in AAI, empirical research examining how knowledge workers interact with these systems in their everyday work remains, to the best of our knowledge, limited. Building on recent calls for empirical research on how workers interact with autonomous AI technologies (Berente et al., 2021; Leonardi, 2025, 2023), and to consider the role of technologies themselves in knowledge work (Newell, 2015), this study investigates how interactions with the distinctive material features of AAI technologies enable and constrain knowledge workers’ ability to achieve work-related goals, using Affordance Theory (Gibson, 1979) as a theoretical lens.

2.3 Theoretical Lens: Affordance Theory

To analyze how knowledge workers’ interaction with AAI’s material properties shapes work practices, we draw on Affordance Theory (Gibson, 1979) as our primary theoretical lens. James J. Gibson initially formulated the Affordance Theory to explain how and why animals utilize their environment in specific ways, based on the proposition that “*what the object affords us is what we normally pay attention to*” (Gibson, 1979, p. 134). The concept has since been adopted in IS research to explain how people interact with new technologies (Hutchby, 2001). In this context, affordances have commonly been understood as action possibilities provided by technologies to goal-oriented actors (Strong et al., 2014; Markus & Silver, 2008; Zammuto et al., 2007). This definition stipulates several critical aspects. First,

affordances should be viewed as relational concepts, in that they exist neither as inherent properties of technology nor as attributes of individual users, but rather emerge from the relationship between users and the technological artefact (Leonardi, 2011). Second, affordances are inherently tied to the goals and purposes of actors, as the same technology may afford different action possibilities to other users depending on their objectives (Leonardi, 2011). Therefore, Affordance Theory provides a middle ground between technological determinism and social constructivism, allowing us to be specific about the technological features of technologies while recognizing that their effects depend on relational dynamics with users and contextual elements (Volkoff & Strong, 2017; Hutchby, 2001; Leonardi & Barley, 2010).

Since the use of Affordance Theory has gained traction in IS research, several explorations and extensions of the theoretical lens have been made. In their study on the implementation and use of electronic health care records (EHR) in hospitals, Strong et al. (2014), for example, made the distinction between affordances as “*potentials for action*” (p.54) and their actualization: “*actions taken by individuals to realize those potentials*” (p.54). This distinction naturally raises ontological questions about whether affordances must be perceived to exist, to which Volkoff and Strong (2013) argue, from a critical realist perspective, that affordances might exist without users being aware of them. However, perception influences the likelihood of their actualization. Their study resulted in a mid-range theory explaining IT-associated organizational change, through what the authors broadly label as the Affordance-Actualization (AA) lens (Volkoff & Strong, 2013). This theoretical lens has since been applied to the study of organizational implementation and individual adoption of technologies such as blockchain (Du et al., 2018) and EHRs (Burton-Jones & Volkoff, 2017).

Another contemporary development is Majchrzak & Markus’ (2012) Technology Affordances and Constraints Theory (TACT), which extends beyond what technologies afford their users also to include explicit attention to their constraints, referring to “*ways in which an individual or organization can be held back from accomplishing a particular goal when using a technology or system*” (p.1). TACT has since been increasingly applied to studies of technology use across a wide range of industries. Yang et al. (2024b), for example, applied TACT to study the use of AI in accounting, revealing both efficiency-related affordances – through enabling the allocation of resources for higher-order tasks – and job-security-related constraints through task displacement.

The Affordance Theory (Gibson, 1979) offers a valuable theoretical lens for studying our research question for several reasons. First, it focuses on interactions and the use of technologies rather than mere adoption or acceptance, aligning with this study's interest in how knowledge workers interact with AAI. Second, it allows us to take a socio-technical perspective by acknowledging the material agency of AAI technologies while avoiding technological determinism, and by recognizing how the interplay among social, contextual, and technological factors actively shapes work practices (Volkoff & Strong, 2017). Third, it centers goal-oriented action, directing attention to what knowledge workers are trying to accomplish rather than simply cataloguing AI features. Finally, the affordance lens offers practical value by revealing how technology features interact with user capabilities to enable or constrain action (Majchrzak & Markus, 2012). This understanding can inform better technology design, more effective implementation strategies, and targeted support for users in actualizing desired affordances (Volkoff & Strong, 2017).

3. Methodology

This chapter aims to guide the reader through the methodological approach adopted. This includes presenting the research approach and design (3.1), the research process and data collection (3.2), the method for data analysis (3.3), and concluding with a discussion of the research quality (3.4).

3.1 Research Design and Approach

3.1.1 Scientific Research Approach

This study adopts a critical realist ontology, in line with previous Affordance Theory scholars (e.g., Gibson, 1979; Volkoff & Strong, 2013), and recognizes that affordances and constraints are relational constructs that emerge from humans' interactions with the material properties of technologies (Leonardi, 2011). This view acknowledges that affordances and constraints are neither inherent properties of the technology nor purely social constructs. Therefore, rejecting the two extreme views of technological determinism (Woodward, 1958) and social constructivism (Orlikowski, 2000). Given that this research examines how knowledge workers perceive and actualize the affordances and constraints through interaction with AAI in their work practices, we employ an interpretivist epistemology to understand this phenomenon. While affordances must not be perceived by users to be actualized, perception may increase their likelihood (Volkoff & Strong, 2017), and understanding this process requires examining how users subjectively experience, make sense of, and respond to the technology's action possibilities.

The research is exploratory and phenomenon-driven (Eisenhardt & Graebner, 2007), reflecting the nascent state of AAI in knowledge work. Rather than testing predefined hypotheses, the study seeks to uncover emergent patterns in how knowledge workers interact with AAI technology and in what influences the actualization of affordances (Strong et al., 2014), thereby supporting the attainment of knowledge workers' work-related goals. In line with the interpretivist perspective, this study adopts a qualitative approach to answer our research question, aligning with Yin's (2003) arguments for its relevance when examining themes in ambiguous and unexplored fields. We posit that the exploratory nature of our research question is best answered with qualitative data (Walsham, 2006).

Further, an abductive research approach was adopted, in which the authors alternated between empirical data, theory, and design choices throughout the study (Bell et al., 2022). Consistent with this abductive logic, the study employed an emergent research design (Lincoln & Guba, 1985), in which the semi-structured interview guide was iteratively refined as new empirical insights emerged. This iterative process allowed for ongoing reflection as new empirical insights emerged from the interviews (Dubois & Gadde, 2002). Through this approach, the Affordance lens was progressively refined to capture the emerging dynamics of knowledge workers' perceptions of action possibilities (i.e., affordances) and the conditions for their actualization.

3.1.2 Research Case

The authors of this study chose to conduct a single case study to uncover the focal phenomenon of knowledge workers' interactions with AAI. Since the affordance lens concerns action possibilities (i.e., affordances) that emerge from interactions between humans and technology, a case study enables deeper insights into this relationship by examining actual user experiences and perceptions when using the focal technology. This method has previously been adopted by IS scholars studying interactions with other technologies through an Affordance lens (e.g., EHR technologies; Strong et al., 2014).

The focal case company, hereinafter referred to as "AgentCo" to maintain confidentiality, is a technology startup licencing an AAI-powered web browser for knowledge workers to use (see section 4.1.1 for a rich case description). While AgentCo serves as the case company for our research, this study focuses on the use of AgentCo's AAI tool rather than investigating AgentCo as a company. Hence, each interview is treated as an individual case unit, as our theoretical interest lies in how individual knowledge workers interact with AAI in their situated contexts. This approach increases variation in our findings and allows us to observe a diverse range of knowledge workers engaging with the same AAI tool (i.e., AgentCo's), thereby capturing individual experiences across different professional contexts, experience levels, and use cases. While we capture how organizational factors (e.g., AI policies, ways of working) shape individual knowledge workers' AAI use, this study does not explicitly focus on organizational dynamics – neither within AgentCo itself nor within the knowledge workers' respective organizations.

The selection of AgentCo as the case company and the decision to study the use of AgentCo's AAI tool were based on two primary considerations. First, one author of this study was employed at the company, providing exclusive access to interview subjects (i.e., users of AgentCo) and deep product knowledge essential for understanding the complexities of users' interactions with AAI's technological features. Second, AgentCo's user base comprises individual knowledge workers across diverse industries and organizational contexts, providing insights into the focal phenomenon of knowledge workers' interactions with AAI. While selecting AgentCo as our case company could introduce certain biases due to one of the researchers being an employee at the company, we deem this risk manageable, as the unit of analysis is individual knowledge workers using AgentCo's AAI tool (i.e., customers) rather than other employees at the case company. However, a further discussion on the mitigation strategies adopted is presented in section 3.4.1.

3.1.3 Research Process

This study employs an abductive research approach using thematic analysis (Braun & Clarke, 2006) and an emergent research design (Lincoln & Guba, 1985), moving back and forth between empirical data collection, data analysis, and engagement with existing theory. The research process is divided into two phases (see Figure 1).




Figure 1. Visual Overview of Research Process

The first research phase is guided by the authors' general interest in the focal phenomenon (i.e., knowledge workers' interaction with AAI), coupled with initial theoretical familiarity

with extant literature on the use of algorithmic technologies in knowledge work. In line with systematic combining principles (Dubois & Gadde, 2002), an interview guide was revised iteratively to probe more deeply into emergent themes and inform the sampling of interviewees. Accordingly, the second phase progressed with an increasing theoretical focus on Affordance theory, specifically on matching empirical themes with emerging theoretical discoveries. Arguments for data collection choices will follow in the next section.

3.2 Data Collection

3.2.1 Phase I

The first phase was informed by our initial interest and theoretical grounding in the focal phenomenon (i.e., knowledge workers' interaction with AAI), leading to the formulation of an open-ended research question and an initial interview guide (see Appendix B) for the first set of interviews. Using purposive sampling (Lincoln & Guba, 1985), these semi-structured interviews were conducted between August and September 2025 with 10 knowledge workers with varying experience using AgentCo's AAI tool (see Table 1). These interviews revealed insights into knowledge workers' use of AAI technologies. Of particular interest were mentions of how the use of AAI either helped or hindered users in their work. These insights informed our choice of theoretical lens (i.e., Affordance Theory). Further, aligning with the propositions of systematic combining (Dubois & Gadde, 2002), these insights informed revisions to the interview guide (see Appendix C), as well as the theoretical sampling for the second phase to provide deeper insights in line with the more developed theoretical lens.

Duration	Role	Industry	Name Code
34 min	Product Manager	Software as a Service (SaaS)	A1
23 min	Solutions Architect	Tech Consulting	A2
30 min	AI Engineer	Cloud Computing	A3
42 min	AI Strategy Consultant	Management Consulting	A4
47 min	Chief Operating Officer (COO)	B2B SaaS	A5
1h 24 min	Operations Lead	B2B SaaS	A6
55 min	Talent Acquisition Specialist	B2B SaaS	A7
35 min	Growth Lead	B2B SaaS	A8
31 min	Product Manager	B2B SaaS	A9
32 min	Strategy Consultant	Management Consulting	A10
32 min	Founder	Real Estate	A11

Table 1. Interview Sample: Phase I

3.2.2 Phase II

The second phase of data collection was conducted between October and November 2025 and comprised 11 additional semi-structured interviews with knowledge workers using AgentCo’s AAI tool (see Table 2). From the first phase, we found that power users (i.e., knowledge workers with extensive experience using the tool) were better positioned than novice users to articulate nuanced accounts of interactions with AAI. Therefore, theoretical sampling (Glaser & Strauss, 1967) was employed to purposefully select participants who were power users. This sampling decision was theoretically motivated by emerging themes on factors that impact the actualization of action possibilities (i.e., affordances) through users’ interactions with AgentCo’s AAI tool, rather than mere perceptions of affordances and constraints. While novice users can identify perceived affordances, power users can speak to the conditions under which affordances are successfully actualized – an emerging theoretical theme in our study. This approach aligns with previous Affordance lens scholars’ emphasis on studying affordance actualization in practice (Strong et al., 2014), which requires observing users who have sufficient experience to move beyond surface-level interactions.

Duration	Role	Industry	Name Code
43 min	Head of Operations	Startup Ecosystem	B1
33 min	Head of Sales	SaaS Fintech	B2
25 min	Head of Operations	SaaS	B3
32 min	Strategy Analyst	Transportation Logistics	B4
42 min	Marketing Consultant	Management Consulting	B5
30 min	Operations Lead	B2B SaaS	B6
37 min	Software Engineer	B2B SaaS	B7
35 min	Data Analytics Consultant	Consulting / E-commerce	B8
40 min	Advisor	Financial Services	B9
60 min	Software Engineer	B2B SaaS	B10
49 min	Digital Strategy Consultant	Consulting / E-commerce	B11

Table 2. Interview Sample: Phase II

As one of the researchers was employed at the company, they had already established contact with these users, allowing us to schedule interviews directly with them. Similar to phase I, these interviews aimed to uncover insights into how the use of AgentCo’s AAI tool enabled or hindered knowledge workers in their professions, but with a deeper theoretical grounding using affordance theory. After completing the 22nd interview across the two phases, no

further interviews were scheduled, as no new theoretical themes or anomalous insights emerged. The final empirical material amounted to 22 interviews (13 hours and 31 minutes of interview recordings) and 154 pages of interview transcriptions. While certain participants (e.g., B10, B5) appear more frequently in quoted material, this reflects their articulateness in describing phenomena rather than unique perspectives. All identified themes were corroborated across multiple interviews; we selected quotes that most clearly exemplified shared patterns.

3.2.3 Interview process

All interviews were conducted online via Google Meet to accommodate interviewees who were not located in Sweden or otherwise busy. Despite video interviews being viewed as a viable means of generating qualitative data and not inherently flawed compared to physical interviews (Deakin & Wakefield, 2013), we still acknowledge some shortcomings in connecting with the interview subject. Therefore, all interviews were led by one researcher, and the other engaged in note-taking and recording non-verbal cues such as body language and facial expressions (Bell et al., 2022). Further, in line with what is commonly referred to as 'typical' ethical practices in qualitative research (Bell & Thorpe, 2013), we deliberately ensured that informed consent was obtained for interview recordings. We deleted video recordings and transcriptions after completing the thesis project.

An interview guide guided the interviews, developed in conjunction with theory and emerging themes identified during the research process and across the two phases of data collection (see Appendices B & C). However, a level of flexibility was also maintained during interviews to accommodate interesting divergences from the intended script. Given the exploratory nature of the research question and the open-ended nature of the interview guide inquiries, this level of flexibility is deemed essential (Alvesson & Sandberg, 2011). All interviews followed the same general sequence: 1) building rapport by informing participants about the purpose of the study and data privacy measures, 2) inquiries for information about background, common workflows, and goals, 3) understanding user motivation for using AAI technologies, and 4) interactions with AAI technologies, including what works well and what does not work well.

We intentionally refrained from using explicit terms associated with our theoretical lens (i.e., the Affordance lens) and instead used conversational terms that were easier to understand.

For example, instead of terms such as 'affordances' and 'constraints', discussions revolved around 'enable' and 'hinder'. This rule of thumb was applied across both interview phases, as conversational language remained beneficial even as theoretical themes began to emerge.

3.3 Data Analysis

Data analysis was conducted in parallel with data collection to allow emerging insights to inform subsequent interviews. Interviews were transcribed using different online transcription tools from Sales Kong, Granola, and Sana. To ensure transcription quality, we reviewed recordings within 48 hours of each interview. Following a thematic analysis approach, which aims to uncover patterns and themes in data rather than imposing preconceived ideas, the data analysis was guided by Braun and Clarke's (2006) six-phase approach. This method was selected for its theoretical flexibility and compatibility with the study's abductive logic, which iterates between theoretical frameworks and empirical data to generate the most plausible explanations for observed phenomena. The data analysis followed a recursive process, with flexible movement between phases as understanding deepened. Further, data from both data collection phases (see section 3.2) were analysed following this process.

After the interviews were transcribed, each was read in its entirety by both researchers to develop an intimate understanding of participants' experiences with AgentCo's AAI. This immersive engagement and familiarization with the data allowed us to identify initial patterns, surprising findings, and potential areas for theoretical exploration.

The initial coding process employed an abductive approach. We began with open coding to identify broad categories and uncover interactions between knowledge workers and the AAI technology, capturing participants' own language and experiences while remaining attentive to patterns that might relate to our theoretical framework. As our understanding deepened, the coding became more focused and selective. Guided by Affordance Theory, we developed codes for users' goals, the specific technological features, and the interactions between them. We then systematically coded for how AgentCo's AAI tool enabled or hindered knowledge workers in their work environment. Throughout this progression, we maintained the flexibility to develop new codes when unexpected or anomalous patterns emerged, ensuring the analysis remained responsive to the data while remaining theoretically informed.

After the initial coding of interviews, we systematically examined both theory-informed and emergent codes to identify broader patterns. Related codes were collated and grouped into themes through an interpretive process that constructed higher-level explanations of how AgentCo’s AAI tool affords or constrains knowledge workers in their pursuit of work-related goals. In addition, a set of contextual factors emerged as anomalous patterns and hence were transferred into themes. Themes were refined, split, merged, or discarded based on their analytical strength and ability to address the research questions. This phase involved iterative movement between data, codes, and the theoretical constructs. Throughout this process, continuous refinement was made to ensure themes were explanatory towards the affordances and constraints related to knowledge workers’ interactions with AgentCo’s AAI tool.

The final analysis is presented as an integrated narrative that weaves together analytical interpretation with quotes, demonstrating how identified themes illuminate the affordances, constraints, and actualization processes of AAI for knowledge workers. The write-up highlights both confirmatory findings (patterns explained by existing theory) and novel insights (patterns that extended or challenged existing theory). Figure 2 presents an excerpt of the analytical coding progression that shows how raw transcript data were turned into first-order codes, grouped into second-order codes, and finally synthesized into overarching themes grounded in Affordance Theory.²




Figure 2. Coding and Data Analysis Excerpt

² see Appendix D for summary of all the affordance and constraints themes found as a result of the data analysis.

3.4 Quality of Research and Ethical Considerations

To allow theoretical and practical contributions to have a substantial impact, a study must exhibit high authenticity and trustworthiness (Merriam & Tisdell, 2016). Drawing on the components of trustworthiness in qualitative research outlined by Lincoln & Guba (1985), this section discusses the quality of this study.

3.4.1 Discussion on Quality

Research credibility is the naturalistic counterpart to internal validity and concerns the degree of congruence between research findings and reality (Lincoln & Guba, 1985). To promote credibility, we first adopted triangulation to corroborate findings by engaging multiple researchers to investigate them (Merriam & Tisdell, 2016). Second, and related, also, to the confirmability of the study – the extent to which findings are grounded in data rather than derived from theoretical inclinations or personal values (Lincoln & Guba, 1985) – one of the researcher’s roles as an employee at AgentCo could, if not handled accordingly, introduce a social desirability bias (Lincoln & Guba, 1985). This dynamic could be dyadic, in which the researcher affiliated with the case company seeks to be perceived in a certain way by users, and vice versa. To combat this potential bias we: (1) explicitly clarified at the start of each interview that the research was academic in nature and that their responses would be anonymized and not shared with AgentCo in an identifiable form, and (2) made sure that the researcher not affiliated with the case company conducted the interviews, while the affiliated researcher took on an observer role and filled in when necessary.

Transferability is instead concerned with the external validity, thereby the applicability of findings to contexts beyond the focal setting of the thesis (i.e., the case organization and the interviewees) (Lincoln & Guba, 1985). To address concerns about the single-case study’s transferability (Bell et al., 2022), we constructed a rich description of AgentCo (see section 4.1.1), thereby enabling peers to effectively judge the transferability of the generated findings (Yin, 2003; Bell et al., 2022). Referring to the consistency of findings and the replicability of the study’s results under similar conditions, dependability is a critical aspect for upholding trustworthiness in qualitative research (Bell et al., 2022). To ensure consistency in findings, we first adopted an auditing approach (Lincoln & Guba, 1985) by providing detailed descriptions of the research process. Second, the support and sequential deadlines offered by our supervisors are argued to serve as an inquiry audit (Lincoln & Guba, 1985).

4. Findings

This chapter presents our empirical findings. Drawing on the Affordance Theory, our findings are organized around themes emerging from interviewees’ descriptions of how AgentCo enables and/or hinders their attainment of work-related goals. Accordingly, we first account for the case context (4.1), offering a deeper description of the case company (i.e., AgentCo), knowledge workers’ goals for using the tool, and AgentCo’s unique technological features. Second, we present five action possibilities (i.e., affordances) enabled at the intersection of users’ goals and AgentCo’s technological features (4.2). Third, an overview of four constraints that prohibit goal attainment is presented (4.3). Lastly, we present three contextual factors that emerged, affecting the likelihood of affordances being actualized (4.4).

4.1 Case Context

4.1.1 The Case Company: AgentCo

AgentCo is a newly founded startup operating in the application layer of AAI. AgentCo integrates the computing power of open-source LLMs to build its own AAI tool available to its users. Founded in 2024, it extends AAI capabilities beyond traditional chatbot interfaces, with an interconnected backend and direct integration into a web browser environment, enabling web-based interaction and workflow automation. AgentCo’s product targets knowledge workers who already do most of their work in a browser, applying AAI within users’ existing work environment. It enables them to deploy autonomous agents capable of browsing the internet, navigating websites, interacting with web interfaces, and executing multi-step workflows across digital platforms – all initiated by natural-language prompts (see Appendix E for a summary on how AgentCo works). A founder of AgentCo describes it as: “*making AI automation accessible to everyone in an intuitive way*”³ (Company website, 2025). According to the founders, their vision centers on fundamentally reshaping how digital work is performed by seamlessly integrating AAI assistance into current work practices, giving knowledge workers a team of AI agents that can automate existing workflows and augment their capabilities, making time for higher-value thinking (Company website, 2025). However, as AgentCo is built on both browser and LLM foundations, users can leverage standard browsing capabilities and conversational AI assistance within the tool. The agentic

³ Quotes and information from the company website has been re-written to maintain anonymity

capabilities are initiated by natural-language instructions from the user, who maintains control over when and how these autonomous agents are deployed.

AgentCo operates in a rapidly evolving market that underwent significant transformation before and during our research period, from August to December 2025. The announcement of a similar web-integrated AAI (i.e., AI browser) by a major AI technology incumbent in early July, alongside similar announcements and product launches from other incumbent actors in the months that followed, indicates growing interest in AAI tools such as AgentCo. As shown in Figure 3, the Google search volume for the terms “AI agents“ and “AI browser“ has increased sharply over the past month, signaling strong public interest in these AAI tools (Google Trends, n.d.).




Figure 3: Google Trends – Search Terms: “AI browser“ and “AI agents“

4.1.2 User Goals

Not surprisingly, the interviews reveal a wide range of context-specific goals among knowledge workers using AgentCo. This is partly explained by the interview sampling decisions and the diverse professional contexts in which the interviewees are situated. However, two themes emerged as particularly salient across most interviews.

First, users are concerned with the quality of the information, the accuracy of the output, and the actions taken by AgentCo. For some, quality was defined as producing accurate output (e.g., from credible sources). In contrast, for others, it was all about creating novel and well-polished information, as B4, a strategy analyst, pointed out:

“I think the quality of the information is the most important thing. Synthesizing the information in a coherent way.” – B4

Second, users were keen on getting more done in less time. It was common in our pool of interviewees to refer to this as “*speed*“, namely, minimising time in executing work to free up time for other tasks. B10, a software engineer, expressed this as:

“Speed of execution is very important. You want to complete as much as possible in as little time as possible. It frees up your capability to do other things.” – B10

We label these two user goals as “ensure quality of work“ and “increase productivity“. The dynamic interaction between these user goals and AgentCo’s technological properties (see section 4.1.3) serves as a set of components that interactively shape the affordances and constraints that emerge throughout the remainder of this findings chapter.

4.1.3 AAI Features

Although nascent, the existing literature on AAI reveals four primary features that distinguish this new set of technologies from previous generations of AI technologies (see section 2.1.2 for a detailed overview). Our findings reveal explicit and implicit mentions of all four features in knowledge workers’ use of AgentCo’s AAI tool. A14, a product manager, for example, mentioned:

“I love the fact that it does the whole breakdown. Like the multiseries agents almost doing research independently.” – A14

This quote highlights the autonomous reasoning and multi-agent coordination features of AAI. Specifically, the words “*breakdown*“ and “*multiseries agents*“ speak to AgentCo’s ability to plan its execution autonomously by breaking a problem into components (i.e., autonomous reasoning) and to autonomously coordinate multiple agents to collectively work on a task (i.e., multi-agent coordination).

Further, B5, a marketing consultant, articulated the persistent memory feature, highlighting that AgentCo can store memories at the agent level, enabling users to leverage specific agentic memories (e.g., instructions, information, and prompts) across different contexts and workflows.

“Because you can have an agent that serves the memory, and then the agent itself goes out and uses different LLMs.” – B5

One of the most persistently mentioned technological features of AgentCo’s AAI tool is the autonomous action-taking (Bornet et al., 2025; Bousetouane, 2025). A10, a strategy consultant, put it simply: *“because it’s read and write”* (A10), highlighting AgentCo’s ability to perform actions rather than only providing text- or image-based output commonly seen in GenAI.

However, a fifth distinctive feature of AgentCo emerges as central across interviews. We label this feature – autonomous web-based interaction – and describe it as: the ability to autonomously access, navigate, and interact with web interfaces, embedding AAI capabilities directly into the users’ browser environment. This native access to the browser environment allows for direct manipulation of graphical user interfaces (GUIs) via button clicks, data extraction, and autonomous web navigation. A10 made an explicit notion of this feature:

“One thing that I really appreciated about [AgentCo] is the fact that it doesn’t need APIs, right? Because it works as an attachment to browsers, it goes from one app to another” – A10

Collectively, these five features – autonomous action-taking, multi-agent coordination, autonomous web-based interaction, autonomous reasoning, and persistent memory – of AgentCo’s AAI tool play a significant role in shaping the action potentials that its users (i.e., knowledge workers) can realize in their work. From hereon out, the terms “AgentCo” and “AgentCo’s AAI” will be used interchangeably to refer to the use of the AAI tool supplied by AgentCo to its users.

4.2 AAI Affordances

The following section outlines the affordances arising from knowledge workers’ interactions with AgentCo’s technological features (see Table 3). These affordances account for the potential actions made possible to knowledge workers, related to their goals (4.1.2), as they interact with the material features of AgentCo (4.1.3).

Foundational Elements of Affordances		Example Quote & Explanation
User Goals	Main AAI Properties	
Affordance 1: Automating workflows		
Increase Productivity Decrease time spent switching between tabs and applications in workflows.	- Autonomous Action-Taking - Autonomous Web-Based interaction	The user does not have to leave the AgentCo interface to take actions on the web or across applications, enabling automated work. <i>“What’s nice about it [AgentCo] is that it’s not hindered by a chat window, right? Because it’s actually using software to go and do things for me.” – A7</i>
Affordance 2: Extracting real-time web-based information		
Ensure Quality of Work Conduct accurate and relevant research.	- Autonomous Action-Taking - Autonomous Web-Based interaction	The user can stay informed by accessing real-time web-based information. <i>“Real-time access is also a game changer, where I feel like it’s [AgentCo] just so much better at finding up-to-date information [than GenAI]” – B5</i>
Affordance 3: Maintaining Contextual Continuity		
Increase Productivity Decrease the time it takes to start and complete tasks.	Persistent Memory	The user does not have to waste time re-prompting AgentCo when switching contexts over time. <i>“I think the memory has helped a lot. When I’ve done research a few days back and it remembers, I think that’s great, and you don’t need to start over, it kind of has context” – B7</i>
Affordance 4: Distributing Specialized Labor		
Ensure Quality of Work Perform high-quality work in multiple domains of expertise across workflows	- Multi-Agent Coordination - Autonomous Reasoning	The user can delegate tasks to specialized agents. <i>“I realized that the best thing for me is to have three different agents. One being the research finder and another being the contextualizer [...] the second agent synthesizes it [the research] and loops it into good themes, and then you have a third agent that actually writes for you” – B1</i>

Table 3. Overview of Affordances

4.2.1 Affordance 1: Automating Workflows

One of the central promises of AAI is that it enables the delegation of concrete actions to agents through autonomous action-taking. AgentCo enables knowledge workers to automate workflows by combining the properties of autonomous action-taking through tool calling via application programming interfaces (APIs) and autonomous web-based interaction via native browser access. Several knowledge workers in our interview sample articulated instances of direct delegation and the intention to delegate tasks to AgentCo to automate specific workflows. For example, a talent acquisition specialist (A7) and a software engineer (B10) mentioned:

“What’s nice about it [AgentCo] is that it’s not hindered by a chat window, right? Because it’s actually using software to go and do things for me.” – A7

“I can do more with less. I can outsource tasks which I think are quite mundane.” – B10

A8, a growth lead, even went so far as automating an outreach workflow to run whilst sleeping, highlighting the automation capabilities offered by AgentCo through both autonomous reasoning and autonomous action-taking

“I did a lot of outreach with massive outreach campaigns, tons and tons of outreach. So I basically had [AgentCo] running 24/7 while I was sleeping.” – A8

Furthermore, a software engineer (B7) used AgentCo to autonomously perform actions across multiple applications in a fundraising workflow, suggesting the possibility of delegating sequential autonomous actions to AAI.

“I’m using [AgentCo] to find VCs in Sweden. You know, find the leads, contact them, and also export them to different tools and apps. For example, if I want to export it to Notion” – B7

Further, many interviewees emphasized the potential to automate workflows directly within AgentCo’s web-based browser. For knowledge workers whose primary work environment is web-based browsers, autonomous action-taking and autonomous web-based interaction offer significant opportunities to automate work. B2, a head of sales in fintech, explained this potential:

“It feels like many of the things I’m doing on Chrome could be automated with AI”. – B2

AgentCo’s integration of AAI into the browser environment enables execution within knowledge workers’ existing workflows, offering significant potential for automation and

time savings. A7, a talent acquisition specialist, articulated the potential impact of such automations, stating: *“This [AgentCo] could cut something like 70% of my workload”* (A7).

4.2.2 Affordance 2: Extracting Real-Time Web-Based Information

Another affordance observed in knowledge workers’ use of AgentCo’s AAI is the extraction of real-time information from web-based sources. Where previous AI applications, such as ChatGPT and Co-Pilot, are bounded by knowledge cutoff points⁴ AgentCo’s unique integration into users’ browsers allows access to real-time web-based information. The core AAI features enabling this affordance are autonomous action-taking and web-based interactions. Several knowledge workers highlighted this affordance. For example, a marketing consultant (B5) and a head of operations (B1) mentioned:

“Real-time access is also a game changer, where I feel like it’s [AgentCo] just so much better at finding up-to-date information [than GenAI]” – B5

“I feel like it is more current research. It [AgentCo] actually went out and looked for the papers [...] it was actually looking for the newest articles and didn’t have the cut-off date that common GenAI tools have”. – B1

The distinction between the capabilities of AAI and GenAI was further articulated by B5, who explained the extended possibilities to autonomously perform research directly within real-time information flows:

“I’ve used it [AgentCo] for scraping social media for clients. And that’s something that I think GenAI can’t do because then they would have to move through social media pages, clicking, analyzing likes and comments [...] real-time information is super important because it changes all the time.” – B5

Beyond accessing current information, AAI’s web-based interaction capabilities enable knowledge workers to extract and synthesize information from diverse online sources at scale. B5, a marketing consultant, emphasized the value of this extraction capability:

⁴ An LLM is trained on data up until a certain point in time. For example, if an LLM was trained on data from before COVID-19, it cannot provide credible answers about specific events that occurred during the pandemic.

“One of my tasks is going through a lot of documents and information, and I think for me, the biggest value add where I’m very impressed is that it [AgentCo] is scraping so much information and extracting it.” – B5

This affordance extends to the discovery of novel information-gathering approaches that knowledge workers may not have previously considered. B3, a head of operations, described how AgentCo autonomously identified an innovative method for lead generation by extracting information from privacy notices:

“It [AgentCo] looked into the privacy notice of a company’s page to find out what system they [the leads] were using. I had not thought about that earlier. So that taught me something new about how you can find leads”
– B3

This example illustrates an action possibility enabled through a combination of autonomous reasoning, autonomous action-taking, and web-based accessibility. For B3, AgentCo was able to both extract real-time data by accessing publicly available web data directly in the browser and display agency by reasoning autonomously on how to execute the task (i.e., finding relevant information about leads) without explicit user prompting.

4.2.3 Affordance 3: Maintaining Contextual Continuity

Many of the knowledge workers interviewed performed a wide range of tasks, regularly requiring them to shift workstreams whilst maintaining accuracy and speed. The persistent memory properties of AgentCo’s AAI proved highly valuable to many knowledge workers, enabling them to store interactions, preferences, and knowledge across contexts. B5, a marketing consultant, described the value this provides:

“I think the memory has helped a lot. When I’ve done research a few days back, and it remembers, I think that’s great, and you don’t need to start over. And it kind of has context”. – B5

B7, a software engineer, further exemplified the persistent memory affordance by contrasting his work with AgentCo with other GenAI tools:

“It [ChatGPT] can’t remember logic. So even if I ask it something, it will forget [...] AgentCo memorizes prompts. For example, my prompts are very long compared to the regular user, but they’re detailed and give me an accurate response every time.” – B7

Storing memory across contexts further allowed this knowledge worker to switch between tasks without having to reintroduce contextual information each time, thereby saving time. A B7 a software engineer, and B8, a data analytics consultant, stated:

“It saves me a lot of time, and at the end of the day, it’s all about time [...] And also, I don’t have to think so much. I can do one feature, for example, work on one project with the team, and then five minutes later jump on the next feature. I can iterate much faster. And actually do much more in less time” – B7

“[AgentCo] can remember things, but in ChatGPT, you need a new prompt for every new case you handle [...] it takes a huge amount of time to do this” – B8

B2, a head of sales, articulated a vision of how this contextual continuity could apply to a group of individuals to generate memories across an organization:

“It should be something that is centralized for the company as a whole. And also so that the tools I use and the LLMs [referring to AgentCo] I ask questions are connected to what the others in the company are asking, so that it continuously learns from what we as a group are doing. Sort of gathers that collective intelligence”. – B2

4.2.4 Affordance 4: Distributing Specialized Labor

Finally, findings show that AAI enables the decomposition of workflows into specialized subtasks, creating processing pipelines that leverage specific distributed expertise. This is in part derived from AAI’s multi-agent coordination feature, which facilitates collaboration between specialized agents through agentic orchestration layers. The ability to coordinate

multiple specialized agents represents a distinctive affordance of AAI that extends beyond what single-agent Gen AI systems provide. B1, a head of operations, describes her creation of multiple specialized agents collaborating across a research workflow:

“I realized that the best thing for me to ensure quality is to have three different agents. One being the research finder and another being the contextualizer [...] the second agent synthesises [the research], and then you have a third agent that actually writes for you.” – B1

This is an example of manual orchestration in which the knowledge worker configures agents’ roles, ensuring each agent specializes in its area of responsibility to deliver high-quality outputs. However, our interviews also uncovered instances of autonomous orchestration. Despite the opacity of agentic orchestration layers, B10, a software engineer, acknowledged:

“You had an agent on top taking the question [...] and then sort of delegated it to a set of different agents who then performed the tasks for me.” – B10

He even continued explaining it as a process similar to how humans break down and delegate tasks to other workers:

“And that felt very orchestrated as if I’m asking a question to someone. Think of it in an organizational context. I ask a question, that person takes in the question, asks something back, and then delegates that to, for lack of a better word, associates or consultants. Then sends out this team who completes the tasks and then comes back with the result in a very structured way” – B10

Ultimately, both the manual (creating several agents) and the autonomous (letting AI orchestrate agent collaboration) aspects of multi-agent coordination enable knowledge workers to complete diverse tasks in a structured way by distributing them to specialized agents.

4.3 AAI Constraints

While the previous section outlined the affordances (i.e., action potentials) at the intersection of AAI and knowledge workers, this section presents findings on how knowledge workers' interactions with AAI's technological properties may also constrain their ability to achieve their objectives (see Table 4).

Foundational Elements of Constraints		Example Quote & Explanation
User Goals	Main AAI Properties	
Constraint 1: Generating Inconsistent Memories		
Ensure Quality of Work Ensure the information produced and research gathered is accurate and relevant	Persistent Memory	Knowledge workers struggled with inaccurate, opaque, and cross-contaminated agentic memories, making it difficult to manage context across workflows. "I want it [AgentCo] to have the learnings from other clients, but I don't want them to have specific client information [...] I think it's hard to create agents that have the correct context without mixing up information from other clients" – B5
Constraint 2: Increasing Idle Time		
Increase Productivity Automate select workflows to increase overall productivity	- Autonomous Action-Taking - Autonomous Web-Based interaction	Knowledge workers experienced trust issues and technical interruptions like CAPTCHA requiring agent supervision and increasing idle work, which limits the productivity potential. "I don't know if it's gonna complete (the task) in a good form or not. [...] With AgentCo, it just has a time, like this is going to take 60 minutes, and then I'm just waiting." – A1
Constraint 3: Struggling With Human-Trivial Tasks		
Increase Productivity Automate simple tasks to increase overall productivity	- Autonomous Action-Taking - Autonomous Web-Based interaction	Knowledge workers are prevented from realizing productivity gains from automation of simple human tasks due to AgentCo's slow and resource-intensive web-based interactions "It [AgentCo] could not find the actual email addresses [...] so it took way longer and used up more credits than it should have" – B6
Constraint 4: Overdelivering on Simple Tasks		
Increase Productivity Get quick answers to simple queries	- Multi-Agent Coordination - Autonomous Reasoning	Knowledge workers are deterred from using AgentCo for simple queries because autonomous reasoning triggers unnecessary multi-step, multi-agent analyses – wasting both time and credits. "It's when you do a simple search and then 12k credits have just been gobbled up. That's what's annoying. Oh my God, I wish I didn't do that." – A11

Table 4. Overview of Constraints

4.3.1 Constraint 1: Generating Inconsistent Memories

While the AAI properties of persistent memory enable some knowledge workers to maintain contextual continuity across workflows (see section 4.2.4), others highlight significant constraints stemming from these properties. Several interviewees mentioned the generation of inconsistent memories (e.g., prompts, contextual information, and preferences) that are not always correct or desired by the knowledge workers. B7, a software engineer, emphasised this constraint:

“A shame in [AgentCo], for example, it’s saving memory, but if I change the context a little bit. There are, like, memories I don’t want to keep” – B7

Knowledge workers in organizational contexts with high demands for information accuracy and quality were particularly prone to experiencing issues with the generation of unintended memories. For example, B5, a marketing consultant, elaborated on this constraint, highlighting the issue of AgentCo spilling information between multiple client workflows:

“When I use the same agent for different clients, she [AgentCo] mixes up the information.” – B5

A few knowledge workers also acknowledged the opacity of agentic memories. A2, a solutions architect, described the difficulty of knowing where information comes from.

“I found in chatting across tabs [...] it remembered something we were talking about from two weeks ago, and I was like, oh, that’s sketchy. Where did that come from?” – A2

The opacity of agents’ memories further made it difficult for knowledge workers to mitigate and manage unwanted memories.

“I want it [AgentCo] to have the learnings from other clients, but I don’t want them to have specific client information [...] I think it’s hard to create agents that have the correct context without mixing up information from other clients” – B5

A digital strategy consultant even opted to turn off the agentic memory feature in AgentCo, stating:

“And when I use [AgentCo], it throws in references just to tell me, look, I can reference memory, but the reference is wrong. It’s out of context. So I actually stopped using memory on [AgentCo] [...] I want it to be as neutral as possible.” – B11

4.3.2 Constraint 2: Increasing Idle Time

As explained in section 4.2.1, the AAI feature of autonomous action-taking enabled multiple knowledge workers to automate workflows by leveraging AgentCo's native browser access to web applications, combined with AAI's ability to act autonomously in that environment. However, several knowledge workers also emphasized a constraint of this action potential: it increases a knowledge worker's idle time when supervising the agent during their action-taking process.

"I'm just sitting and waiting for [AgentCo] to finish tasks, when I could have done something else." – B5

Some knowledge workers found themselves passively supervising AgentCo's actions because they did not fully trust its capabilities.

"I still don't trust everything fully. So it's hard for me to let go of the process completely and do other things." – B5

The action potential of saving time by automating workflows is constrained by individual knowledge workers' need for control.

"I don't know if it's gonna complete [the task] in a good form or not. [...] With AgentCo, it says this will take 60 minutes, and then I'm just waiting." – A1

Further, several knowledge workers noted specific technical interruptions in the agents' environment, requiring manual intervention to keep them from working, ultimately requiring active supervision of AgentCo's autonomous action-taking. For example, CAPTCHA⁵, which are anti-robot mechanisms on the web that require human intervention, force knowledge workers to stay in the loop so the agent can continue operating.

"I feel a little stressed about CAPTCHA coming up that I need to go in and click on to make sure that the workflow continues. [...] It's hard to relax and do other things when you have these processes running at the same time." – B5

Other examples could be log-in pages, checkouts, or moments when the agents need your preferences. In these scenarios, the action potential of saving time by automating workflows

⁵ For example, by indicating which boxes on a grid that contain bicycles is a CAPTCHA that a human must solve to access the content behind the blocker

is constrained by technical interruptions to the autonomous workflow that are part of AgentCo or the environment in which it operates.

“Logging in and just passwords, two-factor authentication, and CAPTCHA, I think, is one of the biggest barriers.” – B4

This imposes unintended time costs on the knowledge worker in the form of idle work time. B5 commented on the value erosion that this dilemma brings:

“It [autonomous action-taking] doesn’t give me much value added, because I still need to look at the page all the time” – B5

B9 further acknowledged that this issue could be exacerbated when AgentCo experiences higher web browsing volume, ultimately discouraging its use.

“Of course, now I can just go into the website and help AgentCo to fix it, but when we do big analyses with many different systems and platforms [...] if it starts to bug [CAPTCHA] every time, then maybe I can do it myself instead of letting the agent do it” – B9

4.3.3 Constraint 3: Struggling With Human-Trivial Tasks

Although many knowledge workers noticed an affordance of automating workflows with autonomous action-taking and autonomous web-based interaction, multiple knowledge workers also acknowledged a constraint derived from these properties: AgentCo struggles to effectively navigate web-based interfaces that are trivial for humans to interact with, which in turn slows task completion and fails to realize the action potential of automating the workflow to increase productivity.

“It said that it can fill out forms for me. And then when I tried it, like, it was actually really, really slow” – A1

“I opened Miro in the browser, and it took a long time to find the sticky note and where to put it [...] in the end I found out that he is looking at a different canvas [...] he did not understand where the middle was.” – B9

Further insights from knowledge workers suggest that, while autonomous action-taking could add value by gathering real-time data, some knowledge workers fail to realize this value because the AgentCo agent navigates web pages more slowly than they do, resulting in higher costs and greater resource (credit) consumption.

“When it [AgentCo] has to move on the website, click, go to the next slide, move down, or do different steps, and not just read a page of information. Sometimes the moving process is very slow. It reads some, then doesn’t understand where to move next.” – B5

“It [AgentCo] could not find the actual email addresses [...] so it took way longer and used up more credits than it should have” – B6

A3 similarly mentioned: *“If it takes too much longer than I would, then it gets a bit annoying”*, highlighting the comparison between the human navigation and that of AgentCo.

4.3.4 Constraint 4: Overdelivering on Simple Tasks

While autonomous reasoning has the potential to augment several capabilities essential to knowledge work (see section 4.3.3) and assist in breaking down tasks for specialized agents (see section 4.3.5), the same feature can also introduce productivity limitations. Several knowledge workers reported instances in which using AgentCo for seemingly simple tasks (e.g., ‘Google-like’ queries and simple executions) triggered complex, time-consuming reasoning sequences. A1, a product manager, recalled an occurrence when searching for information about a vacant job opening using AgentCo:

“[AgentCo] was doing all of these tasks, and it was taking a lot of time, although it’s something that can be found easily. Kind of like phase one research, phase two research [...] and I did not need that kind of deep thinking.” – A1

This instance illustrates over-delivery on a simple task, where AgentCo used autonomous reasoning to break it down into multiple steps, even though no such breakdown was needed. Another knowledge worker, A5, further alluded to the systematic occurrence, pointing out that AgentCo often *“defaults to long, 45-minute research runs,”* which could deter less

committed users. Several interviewees also pointed out the monetary costs of over-delivery on simple tasks:

“So if it’s more simple tasks, it burns credits a lot faster. I don’t wanna use it for simple tasks.” – B5

“It’s when you do a simple search, and then 12k credits have just been gobbled up. That’s what’s annoying. Oh my God, I wish I didn’t do that. Waste of money.” – A6

As indicated by these findings, the unique technological feature of autonomous reasoning could decrease knowledge workers’ productivity by overdelivering on simple tasks, costing time and money through unnecessary credit use.

4.4 Contextual Determinants of Affordance-Actualization

In sections 4.2 and 4.3, the empirical findings introduce a set of affordances and constraints that emerge from the imbrication of AAI properties and knowledge workers’ goals for AAI use. In addition to the action potentials identified, the interviews reveal critical contextual factors that shape knowledge workers’ use of AgentCo and, therefore, affect the likelihood of their actualization.

4.4.1 Individual Characteristics: Ways of Knowing and Need for Control

4.4.1.1 Mental Model Alignment

The interviews revealed distinct patterns in which some knowledge workers benchmark their own problem-solving methods and capabilities against those of AgentCo to determine whether to use AAI to solve a given task. This comparison limits the actualization of affordances derived from AAI features, particularly when knowledge workers face complex tasks. B10, a software engineer, described this deliberation process:

“I don’t think there’s a question of possibility, whether it [AgentCo] can solve those [complex] tasks. I just think that I personally consider them to be harder and trickier. And that’s why I maybe presume that AI is not going to be able to complete them.” – B10

However, B10 also mentioned the need to understand what you are trying to achieve to effectively instruct AgentCo, suggesting that limitations in one's own knowledge of a task may hinder the effective use of AAI:

“It may be that I’m not entirely sure myself how I would solve this problem. [...] In that sense, I’m probably not going to be able to prompt adequately to solve the problem. [...] Easier tickets [tasks] are usually more scoped, so it’s easier to understand what I want, what the input is, and what the expected output is. And it’s easier for me to validate that everything has been conducted correctly.” – B10

The terms “*validate*” and “*correctly*” reflect a desire for control over how tasks are executed and the outcomes generated, based on knowledge workers’ own preferences and knowledge. This indicates that the individual’s need for control, validation, and understanding of the solutions limits the actualization of AgentCo’s action potentials. Several knowledge workers expressed hesitation towards delegating tasks to AgentCo as they struggled to relate the agent’s way of operating to their own and intervene in the technology’s operations:

“When you do something manually, your brain naturally follows every step of the process, [...] it’s harder to intervene, to fix or to backtrack your mistakes when an agent does them” – B4

Other knowledge workers benchmarked AgentCo’s output against previous manual research to “*match what we knew was the truth*” (B3), further indicating evaluation of accuracy against their own knowledge. In a similar vein, B10 noted that the scrutiny of AgentCo’s output might not always be objective:

“The current conceptual model of a system, excluding Agentic AI, is something I understand. However, if it’s [AgentCo] entirely different, then it’s going to be something new to me that I don’t really understand. So that might even bias me to say that, [...] I’m not gonna try and test this because it deviates from how I think the system should look.” – B10

Although the technological features of AAI systems, such as autonomous reasoning, hold the promise of problem-solving and adaptation to complex scenarios, some knowledge workers remain hesitant to use AgentCo in these situations, as its opaque reasoning prevents them from validating the fit between their mental model of problem-solving and the agent's. B2, a head of sales, highlighted this dynamic by noting his desire to understand the technology's underlying mechanisms to use it in more complex workflows.

“In order to have an Agentic AI doing actual complex work for me throughout different workflows, I would feel more comfortable if I understood how it works. [...] I’m not just looking at a black box.” – B2

This “black box” phenomenon directly limits the actualization of the affordance of automating workflows (4.3.1), as the autonomous action potential is not fully realized due to a lack of individual knowledge of the technology.

4.4.1.2 Trust in AI

The lack of trust for AI in general is yet another prevalent factor that has been found to prevent individuals from actualizing the affordances of AgentCo. This lack of trust is partially inherited from previous negative experiences in the broader AI ecosystem, rather than in AAI specifically, and transfers over to their use of AgentCo. B2 explained:

“I think that comes from experience on how wrong answers sometimes can be when you ask an AI something or when you want to do something” – B2

Further, B5 suggests that AI's accuracy limitations are not eliminated by AAI's autonomous reasoning capabilities, leading to a compounded trust deficit when using AgentCo.

“The most frustrating thing is when it gives me inaccurate information. When I’ve used Gen AI or LLMs before, and they give you inaccurate information, you don’t trust it. And then I went into Agentic AI with higher hopes” – B5

B4's experience represents a case where both previous distrust of AI and new distrust for AAI compound into a general feeling of distrust for all tasks requiring more reasoning capabilities:

“I think my impression of AI in general, whether agentic or not, is that it’s very good at doing manual tasks and research, but when it comes to making judgment decisions, it’s not, so I just assume that it will take me longer if I then have to do it myself anyways“ – B4

This lack of trust decreases the likelihood that AgentCo’s action potentials will be actualized, as individuals refrain from exploring and engaging with AAI. However, the converse also holds: the actualization of one action potential fuels continued experimentation. B10 described his process of gradually presenting AgentCo with increasingly complex tasks to avoid disappointment and build trust, explaining:

“That sort of builds up your confidence and your trust in it [AgentCo]. Like, okay, that worked really well. Let’s see what more I can do with it. Let’s see if it can help me to solve even more difficult tasks than that.“ – B10

The same knowledge worker (B10) explained that transparency in autonomous workflows helps increase trust as users can validate the actions by staying in the loop throughout the process:

“You could follow the whole workflow, which helps you to validate, and then you trust it a lot more compared to if you send it into a black box and it outputs, like, 50 emails and like, here you go.“ - B10

However, discrepancies in perceived transparency suggest a level of individuality is present as knowledge workers evaluate the necessary levels of transparency for continued use.

“Even when I went to verify it, I could not find the exact numbers. [...] It might have done some back-end math, studied something, did some sort of calculation, but I want to see what kind of calculation it did“ - A11

4.4.1.3 AAI Literacy

Several knowledge workers expressed a lack of understanding of what to use AgentCo for. This lack of use-case literacy regarding AgentCo's features and capabilities reduces the likelihood of actualizing the technology's action possibilities (i.e., affordances).

“All this agent browser behavior is very new to me, so I'm really trying to figure it out” – A9

This lack of product understanding is highlighted by several knowledge workers requesting specific guidance on how to instruct AgentCo to generate the desired output. B9, B2, B5, and A3 all mention, in different forms, the need for additional assistance to develop “*the best prompts*.” (B9). Some knowledge workers consciously reflected on this literacy gap, demonstrating motivation and a willingness to expand the potential actions they could actualize in the future with AgentCo. B1, for example, said:

“Maybe Agentic AI saves 5% of your time, but if you're ahead of the curve, it can save you 70% of your time in the future. [...] Sure, I'll need to learn prompting, but because this is going to be the future, I don't really mind learning it.” – B1

Several knowledge workers demonstrated absolute AAI illiteracy, using AgentCo exclusively for browser and/or conversational AI capabilities, resulting in the complete underutilization of its AAI features.

“I would say almost all the time I used it [AgentCo] like a one-to-one with Chrome. Obviously, then I didn't use any Agentic AI whatsoever” – B2

4.4.2 Task Characteristics: Stakes, Scale, and Error Tolerance

The nature of tasks emerged as a critical contextual factor determining knowledge workers' willingness to delegate action responsibilities to AgentCo. The possibility of automating workflows (4.2.1) is more likely to be actualized when the task is lower-stakes, such as exploratory and research-related work, as errors carry minimal professional consequences. B5, a marketing consultant, argues:

“I think you always have to have a human in the loop for higher-stakes tasks. I don’t think it’s that necessary if you’re finding LinkedIn leads or stuff like that. There’s no risk there, really” – B5

Conversely, in high-stakes contexts, such as client deliverables and situations in which failure impairs professional reputation, knowledge workers frequently expressed hesitation about using AgentCo to automate workflows. B4, a strategy analyst posited:

“I don’t think I would ever use Agentic AI when the stakes are very high.” – B4

Further, the scale of the task could also influence the willingness to delegate to AgentCo. Multiple knowledge workers noted that AgentCo’s output becomes less controllable and less reliable as scale increases, therefore decreasing their intention to use the tool. B1, a head of operations, for example, said:

“If it’s like a single or a few cases, you can easily fact-check before you do anything. But if it’s a lot of cases, I still wouldn’t use it [AgentCo].” – B1

B11, a digital content manager, noted how scale and stakes often compound:

“I refrain from using [AgentCo] because the output is unreliable. It’s like so much money and people’s lives at stake. If you do something wrong, it can go very, very wrong. [...] On a larger scale, it requires another type of AI system.” – B11

B10, a software engineer, reflected on the burden of responsibility, functioning as a boundary condition for where in the chain of workflows he feels comfortable using AgentCo:

“What happens if I ask you to do something and it fails to execute on that, and something else happens? I guess I’m going to be liable because I outsourced the execution authority to the agent. [...] Now you still have to be like in the beginning, give the input, and then control the output. So it’s like, in the middle, I’m trying to experiment right now.” – B10

He further reflected on the potential value if such constraints are removed:

“I think if that last stage is unlocked, then you start to actually free up a lot more time [...] you can outsource the whole chain.” – B10

While a majority of knowledge workers aligned with B10’s position, preferring to stay in the loop at both ends of a work process (i.e., input and output), particularly for large-scale and high-stakes tasks, some knowledge workers made mentions of exceptions to that rule. For example, A8, a growth lead, automated an entire outreach sequence to AgentCo with no supervision of the output (i.e., actions) generated.

“I did a lot of outreach with massive outreach campaigns, tons and tons of outreach. So I basically had [AgentCo] running 24/7 while I was sleeping.” – A8

Taken together, task stakes and task scale function as boundary conditions that constrain the actualization of automating workflows. However, individual error tolerance can moderate these boundary conditions, as seen by the contrasting views and actions of B10 and A8.

4.4.3 Organizational Context: Restrictive Policies and Ways of Working

Multiple knowledge workers mentioned aspects of the organizational environment that limited their use of AgentCo. First, those situated within larger organizational contexts or who worked for other organizations were generally inclined to express concern about AgentCo’s data security.

“I think one thing for me is that my workplace probably has a little bit of fear, which limits me in the way I can use it” – B5

“I’m an advisor for businesses using AI to grow. Their main question is about data privacy” – A4

“Because I plan to use it [AgentCo] for work, I want to make sure that we’re not going to leak corporate secrets and stuff like that.” – A2

Further, some knowledge workers were strictly prohibited from using AgentCo in their work due to missing or restrictive AI policies rendering certain affordances directly unattainable.

“We don’t use it [AgentCo] because we don’t have an AI policy in our company yet. We’re not allowed to start using different things” – B2

“I’m not allowed to use an Agentic AI” – B10

One knowledge worker pointed to the risk of prompt injection attacks⁶ being the reason for restrictive policies of AAI, specifically within their organization:

“If you drive an agent and hit a website with a malicious prompt hidden in that website, what that can do for data exfiltration is kind of a concern.” - A9

Lastly, organizational ways of working could also affect how and when AAI is deployed. B10, a software engineer, exemplifies the importance of coherence in coding architecture when handing over responsibilities to others in the organization, as a mechanism to deter the use of AAI for coding purposes.

“Once you leave responsibility to someone else, or if I go and work for another company, then the current employer will want everything to be similar because that’s much easier to deal with. And I think that’s sort of a by-product of how software engineering has been done. You have a pattern, you try and follow it. So if I leave that into the hands of an Agentic AI, it might be an entirely different solution than the ones that we typically have.” – B10

The organizational dimension thus operates uniquely compared to other contextual factors: it can create absolute blockages to realizing AgentCo’s affordances at all, or for specific sensitive use cases, rather than merely influencing the degree of actualization possibilities, and it operates primarily outside the individual users’ control.

⁶ LLMs answer to text instructions and lack formal ways of distinguishing instructions from data. Attackers can exploit this vulnerability by embedding hidden commands in prompts (text-based instructions) or external content, tricking AI agents into ignoring original instructions and performing malicious tasks (Debenedetti et al., 2024; Liu et al., 2023).

5. Analysis

This chapter analyzes the interactions between knowledge workers and AAI – synthesizing interviewees’ accounts on how their use of AgentCo enabled or constrained them in their pursuit of work-related goals – through an Affordance Theory lens, developing a conceptual model that explains these interactions (5.1). We then analyze the relational emergence of affordances and constraints (5.1.1) and the conditions enabling their actualization (5.1.2).

5.1 A Conceptual Model of Knowledge Worker–AAI Interactions

Using Gibson’s (1979) Affordance Theory as our theoretical lens, we developed an empirically derived conceptual model showing how knowledge workers interact with AAI to achieve work-related goals (see Figure 4). This model synthesizes our interview findings and reveals the affordances and constraints that emerge when knowledge workers interact with AAI (Majchrzak & Markus, 2012), along with the contextual factors that influence the likelihood of actualizing these affordances (Strong et al., 2014).




Figure 4: AAI-affordances, constraints, and actualization in knowledge work

5.1.1 Affordances and Constraints

Scholars applying the affordance lens to understand human–technology interaction have argued that affordances and constraints are “*relational concepts*” (Majchrzak & Markus, 2012, p. 1) rather than static properties of either technologies or users. Our conceptual model reflects this view by first acknowledging individual knowledge workers’ unique goals (4.1.2) – increasing efficiency and maintaining quality of work – and the key AAI features they interact with to pursue these goals (4.1.3) – autonomous action-taking, persistent memory,

multi-agent coordination, autonomous reasoning, and autonomous web-based interaction. The model then illustrates how these interactions generate both affordances and constraints.

Four key affordances emerged from the interviews, representing action possibilities that enable knowledge workers to accomplish their goals: automating workflows (4.2.1), extracting real-time web-based information (4.2.2), maintaining contextual continuity (4.2.3), and distributing specialized labor (4.2.4). Simultaneously, we identify four constraints that hinder goal accomplishment when using AgentCo: generating inconsistent memories (4.3.1), increasing idle time (4.3.2), struggling with human-trivial tasks (4.3.3), and overdelivering on simple tasks (4.3.4). Positioned at the center of our conceptual model, these affordances and constraints emerge through the relational interactions between knowledge workers' goals and AgentCo's AAI features.

Analysis of findings across affordances (4.2) and constraints (4.3) reveals that all technological features can be perceived as both affording and constraining, depending on knowledge workers' shifting goals. It was further found that affordances and constraints often arise from the interaction of multiple AAI features.

The autonomous action-taking feature, in tandem with autonomous web-based interaction, enables the affordances of automating workflows (4.2.1) and extracting real-time information (4.2.2), allowing knowledge workers to let *"software go and do things for me [them]"* (A7) and *"finding up-to-date information"* (B5). However, for users with productivity-related goals, the same features can be perceived as constraining by increasing idle time (4.4.2) as knowledge workers supervise the agents' operations, and increasing frustration when agents struggle with human-trivial tasks (4.3.3), evidenced by, for example, B5 *"the moving process is very slow"*. Thus, the same AAI features that have the potential to improve quality by retrieving valuable, up-to-date information and increase productivity by freeing knowledge workers from manual work can, in fact, constrain their goals by increasing their idle time and failing to perform actions on human-trivial interfaces.

For users concerned with increasing productivity, the persistent memory feature enabled the affordance of maintaining contextual continuity (4.2.3), allowing them to *"iterate much faster"* (B7) and avoid *"start[ing] over"* (B5) when switching between tasks. However, when the primary focus is on ensuring quality of work, the same feature might be perceived

as constraining, as AgentCo sometimes “*mix[es] up information*” (B5) across workflows (4.3.1), which might be detrimental when working with multiple client-facing workflows. B5’s account introduces a paradoxical dynamic in which the same AAI feature is perceived as both affording and constraining, not only across users but also within the same user, depending on the situated context and goals.

For users seeking to improve the quality of their work, the features of multi-agent coordination and autonomous reasoning in tandem enable the affordance of distributing specialized labor (4.2.4), allowing knowledge workers to explain a more complex task to an AI agent that autonomously “*delegate it to a set of different agents*” and “*then come back with the result in a very structured way*” (B10), similar to the expected actions of delegating work to a human. However, when the primary focus is on productivity and knowledge workers expect trivial answers to straightforward tasks, the same features might come to be perceived as a constraint, as AgentCo might overcomplicate the solution and overdeliver on simple tasks (4.3.4). Hence, the AAI capabilities intended to enable complex task decomposition and completion often become a hindrance for simpler tasks. Even though it improves quality and productivity for some, it occasionally wastes resources and time when knowledge workers want quick answers.

5.1.2 Affordance-Actualization Potential

The right-hand side of the conceptual model recognizes the contextual factors that influence the likelihood that action potentials (i.e., affordances) will be actualized. While the complete actualization process is not systematically traced in this study, these contextual factors shape the conditions under which such actualization is more or less likely.

First, at the individual level, mental models for problem-solving, trust in AI, and AAI literacy (4.4.1), shape the boundaries for when knowledge workers choose to delegate actions to AgentCo and, therefore, whether affordances can be realized in the first place. When knowledge workers are unsure of how to solve a specific problem, they also feel unable to instruct and validate that the actions of an agent have been executed “*correctly*” (B10), and therefore hesitate to use AgentCo for specific complex tasks where autonomous reasoning and multi-agent coordination could, in principle, be most valuable (4.2.2; 4.2.4). Trust deficits inherited from broader AI experiences introduce both barriers to initial use and an increased need for transparency, validation, and control of AAI reasoning. However, trust

also operates dynamically, with the successful use of AAI building confidence iteratively and enabling the progressive delegation of increasingly complex tasks. Lastly, some knowledge workers' AAI illiteracy and insufficient prompting abilities limit the use of AgentCo, ultimately leaving AAI features underutilized or sometimes completely undiscovered.

Second, task characteristics, such as the scale and associated stakes, emerge as determinants of when the action possibilities enabled by autonomous action-taking are welcomed (4.4.2). Most knowledge workers refrain from delegating responsibility to AgentCo for high-stakes, high-scale tasks because they perceive the likelihood and impact of errors to be too high. However, edge cases, such as A8's fully autonomous outreach sequence (4.4.2), suggest that individual error tolerance could moderate the effect of task scale and stakes on willingness to delegate responsibilities to AgentCo.

Third, the organizational context of certain knowledge workers strongly influenced usage and the opportunity to actualize certain affordances (4.4.3). Some organizations operate under restrictive AI regulations, thereby creating absolute blockages for the actualization of specific actions by prohibiting their use by AgentCo. This was seen for B10, who were not allowed to use AgentCo to write code. Interestingly, however, the same knowledge worker acknowledged the potential affordances of automating workflows (4.2.1) and of distributing specialized labor (4.2.4) for other work tasks beyond software development, such as researching and extracting information.

6. Discussion

This study sought to contribute to one of, if not the first, empirical studies on the use of AAI in knowledge work, and thereby close our identified research gaps. The study adopted a case study on knowledge workers' use and interactions with a novel AAI-powered browser. In answering our research question (see below), we address these interactions by examining how the technology affords and constrains actions, as well as how contextual factors shape how actions unfold (6.1). We also discuss our theoretical contributions (6.2), the practical implications of our study (6.3), and provide an overview of the limitations and future research paths (6.4)

RQ: How do knowledge workers interact with AAI to achieve work-related goals?

6.1 Answer to Research Question

Adopting the Affordance Theory (Gibson, 1979) as our theoretical lens, we find that knowledge workers interact with AAI by navigating a complex landscape of affordances and constraints that emerge relationally from their situated goals and AAI's technological features. This interaction generates four key affordances – automating workflows, extracting real-time information, maintaining contextual continuity, and distributing specialized labor – representing action potentials aligned with knowledge workers' goals of increasing productivity and ensuring the quality of work. However, the same technological features can also be perceived as constraints – generating inconsistent memories, increasing idle time, struggling with human-trivial tasks, and overdelivering on simple tasks – that hinder goal attainment.

This duality manifests across all AAI features (see Appendix D). Autonomous action-taking and web-based interaction enable the automation of workflows and the extraction of real-time web-based information, but increase idle time through agent supervision and erode time savings when agents struggle with human-trivial tasks. Persistent memory provides continuity across tasks, saving time for knowledge workers, but can also cause unintended information spillovers between contexts, which is perceived as constraining when information accuracy is essential. Autonomous reasoning and multi-agent coordination enable sophisticated task decomposition and distribution across specialized agents, yet these features can also trigger resource-intensive processes for simple queries, which decreases productivity. Whether these features are perceived as affording or constraining depends on

the alignment of interactions between knowledge workers' situated goals and the technological AAI features.

The likelihood that action potentials (i.e., affordances) are actualized (Strong et al., 2014) is moderated by contextual factors at the individual, task, and organizational levels. Knowledge workers selectively engage with AAI based on their individual characteristics (i.e., knowledge of the task and technology, trust, and alignment with mental models of problem solving); the characteristics of the task (i.e., scale, stakes, and error tolerance); and the organizational context (i.e., policies and organizational ways of working).

6.2 Theoretical Contributions

Our study on the interactions between AAI and knowledge workers through the Affordance lens contributes to several research streams.

First, we extend scholarly understanding of the use of algorithmic technologies in knowledge work by revealing how AAI's unique technological features alter the nature of interactions between humans and technologies that are capable of taking action with limited human supervision (Acharya et al., 2025). We find, as prior research does, that the opacity of algorithmic reasoning poses challenges for interpreting and validating AI output (Faraj et al., 2018). While earlier scholars have demonstrated how knowledge workers interact with and interrogate AI's knowledge claims against their own (Lebovitz et al., 2022; Allen & Choudhury, 2021), our study shows that interrogation practices extend beyond knowledge claims to encompass the AAI's autonomous execution processes as well. Because AAI can autonomously execute tasks on behalf of users, knowledge workers must not only interrogate what agents produce (i.e., knowledge claims) but also how they act.

Accordingly, existing frameworks focused on output-level interactions might be insufficient for explaining interactions between knowledge workers and AAI. Building on Pakarinen & Huising's (2023) finding that accountability shapes algorithmic technology use, we reveal that AAI creates a temporal shift in accountability: knowledge workers must not only interrogate algorithmic outputs but also decide whether to delegate autonomous execution in the first place, as they remain accountable for the actions taken by AAI. This complexity is intensified by the unpredictable and contradictory duality in perceived AAI affordances and

constraints, similar to the “jagged Frontier” described by Dell’Acqua et al. (2023). While AAI can automate workflows, it might simultaneously fail to complete tasks seemingly trivial to a human observer, intensifying the difficulty of predicting where AAI will succeed or fail. Once delegation occurs, autonomous execution proceeds beyond workers’ direct control, which adds increasing tensions to delegation decisions compared to previous generations of AI. Our findings suggest that delegation decisions are shaped by knowledge workers’ perceptions of the autonomous action possibilities afforded by AAI and are moderated by contextual factors, including individual and task characteristics and the organizational context, offering new theoretical insights into what influences the decision to delegate agency for action-taking to AAI technologies (Leonardi, 2025).

Second, we contribute to Affordance Theory (Gibson, 1979) as a lens for understanding human-technology interaction (e.g., Leonardi, 2013; Strong et al., 2014; Majchrzak & Markus, 2012) by revealing situated affordances and constraints for knowledge workers using AAI technologies (i.e., AgentCo). Through our analysis, we identify four affordances and four constraints that emerge from the interactions between knowledge workers and AAI. These findings extend the affordance lens in two ways. First, we demonstrate how knowledge workers interact with AAI features, which are distinct from previous generations of AI, to generate action potentials (i.e., affordances) that are fundamentally agentic rather than informational (e.g., Troicin et al., 2021; Liang et al., 2023). Second, in response to Leonardi’s (2023) call for increased theoretical specification of affordance concepts in studies of AI use, we reveal how contextual factors at the individual, task, and organizational levels shape the potential for affordances to be actualized. While systematic tracking of the full actualization process (Strong et al., 2014) was beyond the scope of this study, our findings offer insights into the conditions under which knowledge workers perceive and engage with autonomous technologies in their work, and when actualization of affordances is more or less likely. Our empirically derived conceptual model integrates user goals, AAI features, affordances and constraints, and contextual factors impacting actualization, offering a framework for understanding how knowledge workers interact with AAI to achieve work-related goals.

Third, our study contributes to one of the first qualitative investigations of AAI use in knowledge work, complementing prior scholars’ predominantly conceptual studies (Acharya et al., 2025; Bornet et al., 2025). By studying knowledge workers’ interactions with AAI during early-stage adoption across diverse professional contexts, we provide rich empirical

insights into a rapidly evolving technological frontier that has received limited empirical attention to date.

6.3 Practical Implications

Our research further contributes to practical implications for three primary stakeholder groups. First, developers of AAI tools in the application layer can use our affordances and constraints to identify areas for improvement to develop functionalities and services that support knowledge workers' attainment of work-related goals. For example, by educating users on the AAI functionalities and possible use cases during onboarding, knowledge workers are more likely to perceive the action possibilities (i.e., affordances) and therefore increase the likelihood of actualization. Further considerations for developers include: adding granular agentic memory controls; implementing smart checkpoints in autonomous action-taking for users who prefer staying in the loop; and making agentic reasoning processes increasingly transparent to increase trust among users who engage in interrogation of autonomous action processes. Second, organizational managers could use our conceptual model to explain how knowledge workers interact with AAI, to inform decisions on how to facilitate the effective use of these emerging technologies. For example, creating sandbox environments and investing in AAI training programs that enable employees to experiment with AAI features, to build greater confidence and knowledge of the potential actions enabled by AAI. Finally, individual knowledge workers could use our study as a roadmap to develop increased AAI literacy as a strategic imperative. Given the rapidly evolving nature of AAI, early adopters will likely enjoy a competitive advantage over their peers.

6.4 Limitations and Future Research

We acknowledge several limitations to our study. First, our single-case design focused on interactions with a single AAI tool (i.e., AgentCo). While this allowed for in-depth exploration and comparison across interviews, our findings may offer limited insights into the affordances of drastically different AAI technologies. Future research could examine interactions across different AAI technologies to identify common patterns and tool-specific dynamics. Second, while we identified contextual factors shaping affordance actualization, we did not explicitly track its outcomes over time (Strong et al., 2014). Longitudinal case studies could offer additional insights into the actualization process and the contextual factors that shape it. Third, our study focused primarily on interactions between individual

knowledge workers and AAI. While we found evidence of organizational factors impacting the affordance actualization process (e.g., AI policies and organizational ways of working), our individual-level analysis did not examine team or organizational dynamics in detail. Future research could investigate how teams coordinate around AAI and how implementation strategies shape the actualization of organizational affordances (e.g., Burton-Jones & Volkoff, 2017; Strong et al., 2014). Finally, as AAI is rapidly evolving, the affordances and constraints identified in our study may be less relevant for future iterations of the technology, necessitating continuous re-examination of human-AAI interaction in knowledge work.

References

- Acharya, D. B., Kuppan, K., & Divya, B. (2025). Agentic AI: Autonomous Intelligence for Complex Goals—A Comprehensive Survey. *IEEE Access*, 13, 18912–18936. <https://doi.org/10.1109/access.2025.3532853>
- Agrawal, A., Gans, J., & Goldfarb, A. (2022). *Power and prediction: The disruptive economics of artificial intelligence*. Harvard Business Press. <https://doi.org/10.1057/s11369-023-00337-z>
- Allen, R., & Choudhury, P. (2021). Algorithm-Augmented work and Domain Experience: the countervailing forces of ability and aversion. *Organization Science*, 33(1), 149–169. <https://doi.org/10.1287/orsc.2021.1554>
- Alvesson, M., & Sandberg, J. (2011). Generating Research Questions through Problematization. *The Academy of Management Review*, 36(2), 247–271. <https://doi.org/10.5465/amr.2009.0188>
- Anthony, C. (2018). To Question or Accept? How Status Differences Influence Responses to New Epistemic Technologies in Knowledge Work. *Academy of Management Review*, 43(4), 661–679. <https://doi.org/10.5465/amr.2016.0334>
- Anthony, C. (2021). When Knowledge Work and Analytical Technologies Collide: The Practices and Consequences of Black Boxing Algorithmic Technologies. *Administrative Science Quarterly*, 66(4), 1173–1212. <https://doi.org/10.1177/00018392211016755>
- Anthony, C., Bechky, B. A., & Fayard, A. (2023). “Collaborating“ with AI: Taking a System View to Explore the Future of Work. *Organization Science*, 34(5), 1672–1694. <https://doi.org/10.1287/orsc.2022.1651>
- Baird, A., & Maruping, L. M. (2021). The Next Generation of Research on IS Use: A Theoretical Framework of Delegation to and from Agentic IS Artifacts. *MIS Quarterly*, 45(1), 315–341. <https://doi.org/10.25300/misq/2021/15882>

Bandi, A., Kongari, B., Naguru, R., Pasnoor, S., & Vilipala, S. V. (2025). The Rise of Agentic AI: A Review of Definitions, Frameworks, Architectures, Applications, Evaluation Metrics, and Challenges. *Future Internet*, 17(9), 404. <https://doi.org/10.3390/fi17090404>.

Bell, E., & Thorpe, R. (2013). *A Very Short, Fairly Interesting and Reasonably Cheap Book about Management Research*. Sage Publications. <https://doi.org/10.4135/9781036208479>

Bell, E., Harley, B., & Bryman, A. (2022). *Business research methods* (6th ed.). Oxford University Press.

Berente, N., Bin Gu, Recker, J., & Santhanam, R. (2021). Managing Artificial Intelligence. *MIS Quarterly*, 45(3), 1433–1450. <https://doi.org/10.25300/MISQ/2021/16274>

Bornet, P., & Wirtz, J. (2025). *Agentic Artificial Intelligence: Harnessing AI Agents to Reinvent Business, Work, and Life*. World Scientific. <https://doi.org/10.1142/14380>

Bousetouane, F. (2025). *Agentic Systems: A Guide to Transforming Industries with Vertical AI Agents*. *arXiv*. <https://doi.org/10.48550/arXiv.2501.00881>

Braun, V., & Clarke, V. (2006). Using Thematic Analysis in Psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>

Brynjolfsson, E., & McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. W.W. Norton & Company.

Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). *Generative AI at work* (NBER Working Paper No. 31161). National Bureau of Economic Research. <https://doi.org/10.3386/w31161>

Burton-Jones, A., & Volkoff, O. (2017). How can we develop contextualized theories of effective use? A demonstration in the context of community-care electronic health records. *Information Systems Research*, 28(3), 468–489. <https://doi.org/10.1287/isre.2017.0702>

Christin, A. (2020). The ethnographer and the algorithm: Beyond the black box. *Theory and Society*, 49(5–6), 897–918. <https://doi.org/10.1007/s11186-020-09411-3>

Coshow, T. (2024, October 1). *Intelligent agents in AI really can work alone. Here's how.* Gartner. <https://www.gartner.com/en/articles/intelligent-agent-in-ai>

Davenport, T. H. (2005). *Thinking for a living: How to get better performances and results from knowledge workers.* Harvard Business School Press.

Deakin, H., & Wakefield, K. (2013). Skype Interviewing: Reflections of Two PhD Researchers. *Qualitative Research*, 14(5), 603-616.
<https://doi.org/10.1177/1468794113488126>

Debenedetti, E., Zhang, J., Balunovic, M., Beurer-Kellner, L., Fischer, M., & Tramèr, F. (2024). Agentdojo: A dynamic environment to evaluate prompt injection attacks and defenses for LLM agents. *Advances in Neural Information Processing Systems*, 37, 82895-82920.
<https://doi.org/10.52202/079017-2636>

Dell'Acqua, F., McFowland, E., III, Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraye, L., Candelon, F., & Lakhani, K. R. (2023). *Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality* (Working Paper No. 24-013). Harvard Business School.
<https://doi.org/10.2139/ssrn.4573321>

Domingos, P. (2015). *The master algorithm: How the quest for the ultimate learning machine will remake our world.* Basic Books.

Dougherty, D., & Dunne, D. D. (2012). Digital Science and Knowledge Boundaries in Complex Innovation. *Organization Science*, 23(5), 1467-1484.
<https://doi.org/10.1287/orsc.1110.0700>

Drucker, P. F. (1957). *Landmarks of tomorrow: A report on the new "post-modern" world.* Harper & Row.

- Du, W., Pan, S. L., Leidner, D. E., & Ying, W. (2018). Affordances, experimentation and actualization of FinTech: A blockchain implementation study. *The Journal of Strategic Information Systems*, 28(1), 50–65. <https://doi.org/10.1016/j.jsis.2018.10.002>
- Dubois, A., & Gadde, L.-E. (2002). Systematic combining: An abductive approach to case research. *Journal of Business Research*, 55(7), 553–560. [https://doi.org/10.1016/S0148-2963\(00\)00195-8](https://doi.org/10.1016/S0148-2963(00)00195-8)
- Eisenhardt, K., & Graebner, M. (2007). Theory building from cases: Opportunities and challenges. *Academy of Management Journal*, 50(1), 25-32. <https://doi.org/10.5465/amj.2007.24160888>
- Faraj, S., Pachidi, S., & Sayegh, K. (2018). Working and organizing in the age of the learning algorithm. *Information and Organization*, 28(1), 62–70. <https://doi.org/10.1016/j.infoandorg.2018.02.005>
- Feuerriegel, S., Hartmann, J., Janiesch, C., & Zschech, P. (2024). Generative AI. *Business & Information Systems Engineering*, 66, 111–126. <https://doi.org/10.1007/s12599-023-00834-7>
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Houghton Mifflin.
- Glaser, B.G., & Strauss, A.L. (1967). *The discovery of grounded theory: Strategies for qualitative research*. Aldine Publishing Company.
- Google Trends. (n.d.). *Interest over time for “AI agents“ and “AI browser“*. Retrieved November 20, 2025, from <https://trends.google.com/trends/explore?q=AI%20browser,AI%20agents&hl=sv>
- Hutchby, I. (2001). Technologies, texts and affordances. *Sociology*, 35(2), 441–456. <https://www.jstor.org/stable/42856294>
- Kaplan, A., & Haenlein, M. (2019). Siri, Siri, in my hand: Who’s the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence. *Business Horizons*, 62(1), 15–25. <https://doi.org/10.1016/j.bushor.2018.08.004>

Kellogg, K. C., Valentine, M. A., & Christin, A. (2019). Algorithms at Work: the new contested terrain of control. *Academy of Management Annals*, 14(1), 366–410.
<https://doi.org/10.5465/annals.2018.0174>

Kim, B., Mehrizi, M. H. R., & Bailey, D. E. (2025). Interlacing Situated and Algorithmic Modes of Knowledge Work: A Workplace Jurisdiction Perspective. *Academy of Management Journal*, 68(5), 907–938. <https://doi.org/10.5465/amj.2023.1286>

Knorr Cetina, K. (1999). *Epistemic cultures: How the sciences make knowledge*. Harvard University Press.

Krakowski, S. (2025). Human-AI agency in the age of generative AI. *Information and Organization*, 35(1), 100560. <https://doi.org/10.1016/j.infoandorg.2025.100560>

Lebovitz, S., Lifshitz-Assaf, H., & Levina, N. (2022). To engage or not to engage with AI for critical judgments: How professionals deal with opacity when using AI for medical diagnosis. *Organization Science*, 33(1), 126–148. <https://doi.org/10.1287/orsc.2021.1549>

Leonardi, P. M. (2011). When Flexible Routines Meet Flexible Technologies: Affordance, Constraint, and the Imbrication of Human and Material Agencies. *MIS Quarterly*, 35(1), 147-167. <https://doi.org/10.2307/23043493>

Leonardi, P.M. (2013). When Does Technology Use Enable Network Change in Organizations? A Comparative Study of Feature Use and Shared Affordances. *MIS Quarterly*, 37, 749-775. <https://doi.org/10.25300/MISQ/2013/37.3.04>

Leonardi, P. M. (2023). Affordances and agency: A clarification and integration of fractured concepts. *MIS Quarterly*, 47(4).
https://misq.umn.edu/misq/article-pdf/47/4/ix/17245/dec-affordances_and_agency.pdf

Leonardi, P. M. (2025). Homo agenticus in the age of agentic AI: Agency loops, power displacement, and the circulation of responsibility. *Information and Organization*, 35(3).
<https://doi.org/10.1016/j.infoandorg.2025.100582>

Leonardi, P. M., & Barley, S. R. (2010). What's under construction here? Social action, materiality, and power in constructivist studies of technology and organizing. *Academy of Management Annals*, 4(1), 1–51. <https://doi.org/10.1080/19416521003654160>

Liang, S., Chang, C., & Huang, C. (2023). Impact of AI Affordance on AI Adoption: The Mediating Role of Self-Monitoring and Uniqueness. *International Journal of Human-Computer Interaction*, 40(20), 6389–6400.
<https://doi.org/10.1080/10447318.2023.2254623>

Lincoln, Y. S., & Guba, E. G. (1985). *Naturalistic inquiry*. Sage Publications.

Liu, Y., Deng, G., Li, Y., Wang, K., Wang, Z., Wang, X., ... & Liu, Y. (2023). Prompt injection attack against llm-integrated applications. arXiv preprint arXiv:2306.05499.
<https://doi.org/10.48550/arXiv.2306.05499>

Loscher, G., & Bader, V. (2022). Augmenting a profession: How data analytics is transforming human resource management. In S. Höllerer, T. Daudigeos, & D. Jancsary (Eds.), *Digital transformation and institutional theory* (pp. 87–110). Emerald Publishing Limited. <https://doi.org/10.1108/S0733-558X20220000083004>

Lyytinen, K., Nickerson, J. V., & King, J. L. (2021). Metahuman systems = humans + machines that learn. *Journal of Information Technology*, 36(4), 427–445.
<https://doi.org/10.1177/0268396220915917>

Majchrzak, A., & Markus, M. L. (2012). Technology Affordances and Constraints in Management Information Systems (MIS). In E. Kessler (Ed.), *Encyclopedia of management theory* (forthcoming). Sage Publications. <https://ssrn.com/abstract=2192196>

Markus, M. L., & Silver, M. S. (2008). A foundation for the study of IT effects: A new look at DeSanctis and Poole's concepts of structural features and spirit. *Journal of the Association for Information Systems*, 9(10), 609–632. <https://doi.org/10.17705/1jais.00176>

McCarthy, J. (2007). *What is artificial intelligence?* Stanford University.
<https://www-formal.stanford.edu/jmc/whatisai.pdf>

- Merriam, S. B., & Tisdell, E. J. (2016). *Qualitative research: a guide to design and implementation* (4th ed.). John Wiley & Sons.
- Mukherjee, S. (2017, April 3). *A.I. versus M.D.: What happens when diagnosis is automated?* The New Yorker. <https://www.newyorker.com/magazine/2017/04/03/ai-versus-md>
- Newell, S. (2015). Managing Knowledge and Managing Knowledge Work: What we know and what the Future holds. *Journal of Information Technology*, 30(1), 1–17.
<https://doi.org/10.1057/jit.2014.12>
- Orlikowski, W. J. (2000). Using technology and constituting structures: A practice lens for studying technology in organizations. *Organization Science*, 11(4), 404–428.
<https://doi.org/10.1287/orsc.11.4.404.14600>
- Pachidi, S., Berends, H., Faraj, S., & Huysman, M. (2020). Make way for the algorithms: Symbolic actions and change in a regime of knowing. *Organization Science*, 32(1), 18–41.
<https://doi.org/10.1287/orsc.2020.1377>
- Pakarinen, P., & Huising, R. (2023). Relational Expertise: What Machines Can't Know. *Journal of Management Studies*, 62(5), 2053–2082. <https://doi.org/10.1111/joms.12915>
- Pettersen, L. (2019). Why artificial intelligence will not outsmart complex knowledge work. *Work, Employment and Society*, 33(6), 1058–1067.
<https://doi.org/10.1177/0950017018817489>
- Prawat, R. S. (1996). Constructivisms, modern and postmodern. *Educational Psychologist*, 31(3), 215–225. <https://doi.org/10.1207/s15326985ep3103>
- Shavit, Y., Agarwal, S., Brundage, M., Adler, S., O'Keefe, C., Campbell, R., ... Robinson, D. G. (2023). *Practices for governing agentic AI systems*. OpenAI.
<https://cdn.openai.com/papers/practices-for-governing-agentic-ai-systems.pdf>

Strich, F., Mayer, A.-S., & Fiedler, M. (2021). What do I do in a world of artificial intelligence? Investigating the impact of substitutive decision-making AI systems on employees' professional role identity. *Journal of the Association for Information Systems*, 22(2), 304–324. <https://doi.org/10.17705/1jais.00663>

Strong, D., Volkoff, O., Johnson, S., Pelletier, L., Tulu, B., Bar-On, I., Trudel, J., & Garber, L. (2014). A Theory of Organization-EHR Affordance Actualization. *Journal of the Association for Information Systems*, 15(2), 53–85. <https://doi.org/10.17705/1jais.00353>

Sukharevsky, A., Kerr, D., Hjartar, K., Hämäläinen, L., Bout, S., Di Leo, V., & Dagorret, G. (2025). *Seizing the agentic AI advantage*. McKinsey & Company. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/seizing-the-agentic-ai-advantage>

Talbert, M. (2025, April 17). *How work about work gets in the way of real work*. Asana. <https://asana.com/sv/resources/why-work-about-work-is-bad>

Trocin, C., Hovland, I. V., Mikalef, P., & Dremel, C. (2021). How Artificial Intelligence affords digital innovation: A cross-case analysis of Scandinavian companies. *Technological Forecasting and Social Change*, 173, Article 121081. <https://doi.org/10.1016/j.techfore.2021.121081>

Van Den Broek, E. (2025). Unpacking AI at work: Data work, knowledge work, and values work. *Information and Organization*, 35(3), 100584. <https://doi.org/10.1016/j.infoandorg.2025.100584>

Van Den Broek, E., Sergeeva, A., & Vrije, M. H. (2021). When the Machine Meets the Expert: An Ethnography of Developing AI for Hiring. *MIS Quarterly*, 45(3), 1557–1580. <https://doi.org/10.25300/misq/2021/16559>

Volkoff, O., & Strong, D. M. (2013). Critical Realism and Affordances: Theorizing IT-Associated Organizational Change Processes. *MIS Quarterly*, 37(3), 819–834. <https://doi.org/10.25300/MISQ/2013/37.3.07>

Volkoff, O., & Strong, D. M. (2017). Affordance theory and how to use it in IS research. In R. D. Galliers & M.-K. Stein (Eds.), *The Routledge companion to management information systems* (pp. 232–245). Routledge. <https://doi.org/10.4324/9781315619361-18>

Waardenburg, L., Huysman, M., & Sergeeva, A. V. (2022). In the land of the blind, the one-eyed man is king: Knowledge brokerage in the age of learning algorithms. *Organization Science*, 33(1), 59–82. <https://doi.org/10.1287/orsc.2021.1544>

Walsham, G. (2006). Doing interpretive research. *European Journal of Information Systems*, 15(3), 320–330. <https://doi.org/10.1057/palgrave.ejis.3000589>

Woodward, J. (1958). *Management and technology*. H.M. Stationery Office.

Yang, J., Amrollahi, A., & Marrone, M. (2024). Harnessing the Potential of Artificial Intelligence: Affordances, Constraints, and Strategic Implications for Professional Services. *The Journal of Strategic Information Systems*, 33(4), 101864. <https://doi.org/10.1016/j.jsis.2024.101864>

Yee, L., Chui, M., Roberts, R., & Xu, S. (2024). *Why agents are the next frontier of generative AI*. McKinsey & Company. <https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/why-agents-are-the-next-frontier-of-generative-ai>

Yin, R. (2003). *Case study research: Design and methods*. Sage Publications.

Zammuto, R. F., Griffith, T. L., Majchrzak, A., Dougherty, D. J., & Faraj, S. (2007). Information technology and the changing fabric of organization. *Organization Science*, 18(5), 749–762. <https://doi.org/10.1287/orsc.1070.0307>

Appendices

Appendix A. Comparison of traditional AI and agentic AI capabilities

Category	Traditional AI	Agentic AI
Primary Purpose	Task-specific automation	Goal-oriented autonomy with complex, multi-step execution
Autonomy Level	High human intervention – requires predefined parameters	Low human intervention – operates independently over time
Memory & Context	Limited memory; no persistent memory across sessions	Persistent memory (remembers past interactions, adjusts plans accordingly)
Learning & Adaptation	Primarily supervised learning; static until retrained by developers	Learns from interactions and refines behaviour; uses reinforcement and self-supervised learning
Decision-Making	Data-driven, static rules; rule-based or supervised learning	Autonomous, contextual reasoning with dynamic adaptation to changing environments
Task Complexity	Handles single, well-defined tasks with clear parameters	Manages complex, multi-step tasks with ambiguous requirements

Adapted from Acharya et al. (2025) and Bornet et al. (2025).

Appendix B. Interview Guide: Round 1

Interview Guide [Round 1]

Info [2min]

This interview is part of a research project and it's about understanding how you use Agentic AI. Getting optimal use out of agentic AI tools is not always easy, and there is often a lot of trial and error on the way to reaching optimal ways of working with AI agents – we would like to understand how you experience this, specifically your experience with using AgentCo.

The interview is completely anonymous and will be recorded with your consent (to be deleted after the research paper is written).

Background Discovery [5-10min]

Understand their background, initial expectations, and early motivations for trying Agentic AI.

- What do you do? Tell me about your work.
- What industry are you in? What kind of company?
- How did you hear about us? What made you try AgentCo?
- What have you been using it for so far?
- What were you hoping AgentCo would help you with when you first tried it?
- Did you have any expectations or concerns before starting to use an agentic AI tool? What did you expect AgentCo to do differently than the other tools?

Understanding Their Workflow [5-10min]

Deeply understand current work processes and what workflows they would integrate AgentCo into, as well as the underlying jobs-to-be-done and pain points.

- Walk me through a common work process or task that you do often.
- How do you currently handle your most time-consuming tasks?
- What tools are you using now for this?
- What's your biggest time sink right now?
- Where do you see the biggest friction in your daily workflows?

Value Exploration & Feedback [5-10min]

Explore their concrete experiences using AgentCo, including successes and breakdowns.

- What worked well with AgentCo? What didn't work?
- Where did it get stuck? What issues did you run into?
- How long did it take? How many credits did it use?
- Run me through a specific task you used AgentCo for and how it went.
- What did you do to mitigate issues or work around them?

Exploring Use Cases [5-10min]

Understand how they imagine the tool evolving within their workflows.

- What else could you see using this for?
- What would be most valuable for you?
- What's the biggest time-saver you could imagine?
- How do you see agentic AI fitting into your current ways of working going forward? What would it take to completely implement AgentCo into your daily workflows?
- Has your way of using AgentCo changed since you first started?

Ending [3-5min]

“Any final reflections before we wrap up?”

Appendix C. Interview Guide: Round 2

Interview Guide [Round 2]

Info [2min]

This interview is part of a research project and it's about understanding how knowledge workers interact with Agentic AI. Getting optimal use out of agentic AI tools is not always easy, and there is often a lot of trial and error on the way to reaching optimal ways of working with AI agents – we would like to understand how you experience this, specifically your experience with using AgentCo. The interview is completely anonymous and will be recorded with your consent (to be deleted after the research paper is written).

Background [3–5min]

Understand their background and previous experience with AI agents.

- Briefly describe your profession, role, and the industry you operate in.
- How much experience do you have with Agentic AI tools (AgentCo included)?

Work & User Goals [5–10min]

We would like to understand how you work, your goals and what is important to reach your goals.

- Walk me through what you do on a daily basis.
 - Common tasks? Workflows? Tools/actors involved?
- What does good look like in your daily work?
 - What are your goals? (speed, quality, accuracy, etc.)
- What is the most difficult part of your job?
- Are there parts of your workflow you see most potential to optimize or automate?

Motivation: Why Agentic AI? [3–5min]

- Why did you choose to work with Agentic AI?
- What expectations and beliefs did you have at the start?
- What did you want the tool to do / not do? What did you want to achieve?
- Any fears? Why? Examples.

User Interactions with Agentic AI (AgentCo) [15–35min]

Understand how you interact with AgentCo and how it fits into your work. To understand the imbrication of AAI features and user goals that leads to affordances and constraints. Taking it one step further to get an understanding if their use of AgentCo leads to any improvement towards their work-related goals.

- Please tell me about the last or a recent time you used AgentCo.
 - What did you want to accomplish? How did it go?
- Walk me through a moment when AgentCo helped you a lot.
 - Why did it work so well?
 - What did you do to ensure this?
 - Describe in detail how you collaborated with AgentCo.
- Walk me through a moment when AgentCo didn't work as you wanted.
 - Frictions? Frustration? Examples? Any learnings?
 - What did you do to mitigate this / workarounds?
 - What was lacking from you, your environment or the technology?
- Do you have any strategies to make better use of the AAI or reduce issues?
- In comparison with other AI tools that you've used, what stands out with AAI and AgentCo for you? Elaborate
- Has your way of working with AgentCo changed since the first time you used it? If so, how has it evolved?
- Are there aspects of the technology (e.g., autonomy, black-boxed nature, hallucinations) that make it more challenging than regular tools or doing it manually?
- What value have you seen from this service in the end, personally or for your workplace?
- In what way has it contributed or not to the goals you mentioned in the start?

Ending & Final Thoughts [3min]

Any final reflections or examples you think are important for us to understand?

Appendix D. Summary of Affordances and Constraints: Grouped by AAI Feature

Main AAI Properties	Affordances	Constraints
1. Autonomous Action-Taking AAI can act autonomously by using tools through API connectors	(P) Automating Workflows (Q) Extracting Real-Time Web-Based Information	(P) Increasing Idle Time (P) Struggling With Human-Trivial Tasks
2. Autonomous Web-Based interaction AAI can perform real-time action directly within the users' browser environment	(P) Automating Workflows (Q) Extracting Real-Time Web-Based Information	(P) Increasing Idle Time (P) Struggling With Human-Trivial Tasks
3. Persistent Memory AAI can aggregate memories (e.g., instructions and interactions) across different contexts (sessions and agents).	(P) Maintaining Contextual Continuity	(Q) Generating Inconsistent Memories
4. Multi-Agent Coordination AAI can assign responsibilities to specialized sub-agents	(Q) Distributing Specialized Labor	(P) Overdelivering on Simple Tasks
5. Autonomous Reasoning AAI can break a task into relevant sub-actions it deems necessary to solve it.	(Q) Distributing Specialized Labor	(P) Overdelivering on Simple Tasks

Note: "Q" denotes the user goal "*ensure quality of work*", and "P" denotes the user goal "*increase productivity*"

Appendix E. How AgentCo Works: Answered by an AgentCo Agent

Prompt	“Can you tell me about how you are built?”
Output	“Hey! I'd be happy to share what I can about how I work within [AgentCo]. I'm [name of specific agent], an AI assistant that lives inside the [AgentCo] Browser. Here's what makes me useful for you: What I can do: • Browse and interact with websites - I can navigate sites, click buttons, fill forms, extract information, and automate tasks just like you would • Research deeply - I can deploy several sub-agents online to search databases, websites, and any source online, and then compile the information. • Analyze documents - I can read and analyze PDFs, papers, and other attachments you share with me • Use integrations - I can connect with apps like Gmail, Google Sheets, Notion, Google Docs, and others to help organize your work • Remember context - I learn your preferences and study topics over time to be more helpful in future conversations • Multi-task - For large projects, I can work on multiple tasks in parallel to get things done faster. How I work: • I operate from your browser and can access the same websites you can (assuming you're logged in) • I use specialized agents to browse and research - think of them as focused helpers I send out to complete specific tasks • I'm rate-limited to normal human browsing speed, so websites see my actions as regular user activity

Appendix F. Disclosure of AI Use

Throughout this research, we employed several GenAI-powered tools to enhance efficiency while maintaining academic integrity. This appendix details these tools, their applications, and the verification strategies we implemented.

AI Transcription Tools

We used three AI-based transcription services – Granola, Sales Kong, and Sana Labs – to transcribe interview recordings and generate summaries. The summarization and search functionalities helped us efficiently recall interview content during analysis, particularly when revisiting specific themes across multiple interviews. However, we conducted our data analysis and coding without AI assistance, using these tools only to validate and challenge our identified themes after completing the data analysis. Automated transcription carries risks of inaccuracies, including misinterpreted words or missed nuances. To mitigate this, we reviewed transcripts alongside original audio recordings within 48 hours to identify and correct transcription errors, ensuring that our analysis accurately reflected participants' responses. Additionally, both authors participated in all interviews and took manual notes (with one researcher bearing primary responsibility for note-taking). All recordings and transcripts were handled in accordance with ethical guidelines and will be deleted upon thesis completion.

Example prompts: *“What are the main takeaways from this interview?” “What are the key themes across all of these interviews?” “Find me the quote about [a specific mention] in this interview.”*

AI Writing Aid Tools

We used Grammarly throughout the writing process to identify grammatical errors and suggest alternative sentence structures. The tool functioned as a quality assurance mechanism, helping detect errors that might be overlooked during manual proofreading. Given that large language models can occasionally produce erroneous suggestions, we critically evaluated all recommendations within the full context of each sentence and paragraph. We accepted only suggestions that preserved the intended meaning and enhanced clarity. The tool was not used to generate content but solely to refine our existing text.

General AI and Agentic AI Tools

We used ChatGPT (5.1), Claude (Sonnet 4.5), and AgentCo (Claude 4.5 and other models in an agentic orchestration embedded on the web within a browser) in a limited and carefully controlled manner for: (1) feedback on draft passages, (2) suggesting alternative formulations for clarity, and (3) assisting in web searches to identify potentially relevant academic sources during early literature review.

These tools were not used to generate complete sentences or sections of the thesis. Instead, they served as consultative aids to refine our content and expedite initial literature identification. All AI-suggested alternative formulations were critically reviewed and adapted to maintain the authors' academic integrity and voice. To emphasise, AI tools were not used to generate text and meaning beyond what had already been outlined and written solely by the authors. A significant risk when using generative AI for literature searches is the potential for fabricated sources – a phenomenon known as “hallucination.” To address this issue when using AI to assist in broad web searches, we independently verified all AI-suggested sources through academic databases and search engines. Only sources we confirmed existed and demonstrated relevance were included in our literature review. Furthermore, we accessed and read all cited works in their original form to ensure accurate representation.

Example prompt (Claude/ChatGPT): *“Without altering the originality and emphasis of this text, suggest improvements to grammar, flow, and readability for me to review.”*

Example prompt (AgentCo): *“Start agents on the web to find top-rated academic journals on [a specific topic]”*

Ethical Considerations

Principles of transparency, verification, and human oversight guided our use of AI tools. While these tools enhanced efficiency, all substantive intellectual work – research design, data collection, analysis, interpretation, and writing – was conducted by us as researchers. AI tools served supplementary functions and did not replace critical thinking or scholarly judgment. We treated all AI outputs as suggestions requiring human verification rather than authoritative information, ensuring our research integrity and originality remained intact.

Generative AI Sources

AgentCo. [N/A⁷] (confidential as it is the case company of this study)

Anthropic. (2025). [Sonnet 4.5] Claude. <https://claude.ai/>

Grammarly. (2025). [N/A] Grammarly. <https://grammarly.com/>

Granola. (2025). [N/A] Granola. <https://granola.ai/>

OpenAI. (2025). [GPT-5.1 Thinking] ChatGPT. <https://openai.com/>

SalesKong. (2025). [OpenAI] SalesKong. <https://saleskong.com/>

Sana Labs. (2025). [N/A] Sana Labs. <https://sanalabs.com/>

⁷ N/A denotes the tool is either auto-routing queries to multiple LLM-models to generate output, or that the LLM-model used is undisclosed.