Goals as First-Class Abstractions in Human-AI Collaboration

LEV TANKELEVITCH, Microsoft Research, United Kingdom
SEAN RINTEL, Microsoft Research, United Kingdom

As AI assumes more of the material production in knowledge work, human effort shifts toward planning, orchestration, and evaluation, all of which revolves around goals. Yet goals remain poorly represented in knowledge work tools and workflows: implicit, unexpressed, or confused with outputs. Beyond their importance for human work, clear goals are fundamental to human-AI communication and collaboration. We review research establishing the value of explicit goals, show through a review of commercial tools that existing ecosystems support goal tracking but not goal articulation, alignment, or contextual use, and use meetings as a proving ground demonstrating that upstream goal articulation produces disproportionate downstream value for both humans and AI agents. We argue that goals should be encoded as first-class abstractions that drive human-AI collaborative workflows and that generative AI's natural-language capabilities make this a uniquely opportune moment. We outline six design requirements for goal-oriented human-AI collaborative systems.

CCS Concepts: • **Human-centered computing** → **Computer supported cooperative work**; *HCI design and evaluation methods*; • **Computing methodologies** → *Multi-agent systems.*

Additional Key Words and Phrases: goals, goal articulation, human-AI collaboration, knowledge work, agent orchestration, generative AI, goal-setting theory

ACM Reference Format:
Lev Tankelevitch and Sean Rintel. 2026. Goals as First-Class Abstractions in Human-AI Collaboration. In *Proceedings of AutomationXP26 Workshop of the 2026 CHI Conference on Human Factors in Computing Systems, April 14, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 7 pages.

1 Introduction

As AI takes over the material production of knowledge work, humans' roles are shifting toward planning, orchestration, and evaluation: deciding what work to do, who should do it, and whether it's on track [29, 33, 40, 46]. Human workers are becoming managers of hybrid human-AI teams organised around *goals*. We define goals broadly as desired outcomes at any altitude, from task-level to strategic, hierarchically related, and often dynamically evolving [14]. Goals are the organising backbone of knowledge work, allowing teams to prioritise, coordinate, persist, and track progress [1, 14, 16, 24, 25, 53]. They are equally important for human-AI collaboration [4, 20]. As Peter Drucker emphasised, "*working on the right things is what makes knowledge work effective*" [14].[1] Despite this, even human teams struggle with a core challenge: articulating and coordinating around goals [6, 38, 42]. The complexity introduced by AI's accelerating productive capacity and multi-agent workflows threatens to exacerbate this [7]. Thinking clearly about goals will not

[1]Although, as emphasised by Drucker [14], goal articulation and clarity is *essential* to knowledge work given the relatively abstract nature of its materials and output, the organising function of goals and the design requirements that we outline here are also relevant for broader forms of work.

Authors' Contact Information: Lev Tankelevitch, Microsoft Research, Cambridge, United Kingdom, lev.tankelevitch@microsoft.com; Sean Rintel, Microsoft Research, Cambridge, United Kingdom, serintel@microsoft.com.

be automated away [4, 20]. Rather, to enable effective human-AI collaboration, we need to rapidly augment people's ability to articulate and coordinate around goals.

Despite the centrality of goals in collaborative knowledge work [24], they remain poorly represented in its artefacts, workflows, and communication [6, 38]. Even popular frameworks like Objectives and Key Results (OKRs) [12] face persistent implementation challenges [6, 19, 42]. Goals often exist entirely in people's minds, either unformed or unexpressed, often leading to a confusion of *outputs* for *outcomes* [12, 38, 53]. Meetings, decks, and shipped features are treated as ends rather than means. If goals are as effective as decades of research and practice suggest, then we must do better. OKR planning sessions and task management tools shouldn't be the only places where people and AI agents come into close contact with their goals.

All the above will become more important and more complex in a world where human-AI teams collaborate on knowledge workflows [45, 48]. At a basic level, clearly expressed goals—e.g., via prompts to Generative AI systems—are core to the formation of *common ground* between humans and AI systems, ensuring that parties are aligned on tasks and responsibilities [4, 10, 20]. Yet, as research on Generative AI has found, effective articulation of goals via prompting is challenging for users, with systems insufficiently supporting this process [46]. Moreover, if AI shifts human labour to more managerial roles, as has been anticipated [14] and observed in prior decades of technological transformation [17, 39], many more workers will need to be effective at goal articulation, prioritisation, coordination, and tracking [21, 33, 46]. Delegation and coordination among human-AI teams will also need to account for the scale differences in operation and communication between human and AI agents [4, 40, 48, 52], as well as changes in workflows: for example, workers with access to AI have started to cross disciplinary boundaries [43] and blur the line between planning, ideation, and production processes [21]. Given these challenges, we argue that human-AI collaborative systems should:

(1) Support the articulation of goals (i.e., desired outcomes)
(2) Encode goals as explicit abstractions
(3) Use these abstractions to drive goal-oriented experiences to shape people's (and AI agents') thinking, communication, and workflows towards better outcomes, rather than more outputs.

2 The Value of Explicit Goals

Decades of research on goal-setting theory [24, 26] and industry practice around OKRs [12, 53] establish four functions that explicit goals serve, each amplified by AI's arrival in the workflow:

(1) Goals **direct attention**: At the individual level, goal articulation serves a metacognitive function, increasing self-awareness (e.g., of trade-offs) and enabling task decomposition and planning [27, 37, 46]. In collaborative settings, goal articulation surfaces critical disagreements and divergent assumptions that might otherwise remain implicit within teams [14, 53]. The need for focus and prioritisation only increases in an AI-first world. As AI cheapens the *production* of knowledge work, the cost to human *attention* increases [30, 35], and the bottleneck therefore shifts to being able to prioritise the right kind of work [33, 34].
(2) Goals **enable flexible coordination**: shared goals let collaborators, human or AI, interpret and infer others' behaviour without tracking every action, reducing communication overhead [24, 41, 48]. As agents proliferate, such coordination becomes essential to maintain alignment [4, 20]. This flexibility increases the robustness of human-agent workflows to the *situatedness* of collaborative work [44]: action is contextual and plans may change, yet work remains goal-oriented. Human-AI teaming also implies that expressed goals will need to be interpretable and actionable by both humans and agents.

(3) Goals also **drive persistence** by providing clear endpoints and enabling self-regulation against an external standard [26, 31]. For AI, they serve as objective functions that shape optimisation and resource allocation. A recent example is Claude Code's *Ralph Wiggum* plugin, which enables a continuous self-referential feedback loop that drives persistent iteration until all requirements are met [2].
(4) Finally, goals **enable evaluation** [18, 53]: they provide benchmarks for success and iterative improvement. In human-AI collaboration, this evaluative function is bidirectional: goals let humans assess AI outputs against intended outcomes [4, 5] and let AI provide goal-referenced feedback to human collaborators [8].

Given these functions, the challenge is making goals available not just in quarterly planning, but continuously, at the point of work, for both humans and agents.

3 Current Goal Tools Provide Tracking Without Thinking

A review of current goal tools (task managers, project platforms, and dedicated goal-setting products) reveals that a handful (Asana[2], ClickUp[3], Atlassian Goals[4]) offer explicit goal objects, hierarchical alignment, and status rollups. But these are foundations for *tracking* goals, not for *thinking with* them. The gaps are revealing. No tool supports goal *articulation*: the reflective process of working out what you are actually trying to achieve. Where AI assistance exists, it is syntactic ("rewrite this goal more clearly") rather than substantive. Goal features are confined to individual products. No productivity suite threads goal awareness across its own applications, let alone the broader tool ecosystem, leaving goals stranded in silos. Collaborative goal-setting is essentially unsupported: no tool surfaces misalignment or implicit goals, and none supports structured negotiation toward shared goals. Goals are invisible at the point of work: no system provides situated goal cues during workflows, and no tool supports goal-guided agent orchestration. Tracking itself is shallow: no tool captures the *why* behind a goal, completed goals vanish, and there is no predictive capability. Existing tools assume goals are set once and occasionally reviewed. They are not built for a world where goals are living abstractions that shape how people and agents coordinate.

4 Example: Meetings Goals and Downstream Value

Meetings offer a compelling test case: ubiquitous, consequential, and chronically under-served by goal support. Microsoft's 2023 Work Trend Index identified inefficient meetings and a lack of clear goals as the top two obstacles to workplace productivity [28]. Research confirms that unclear goals are a critical factor in meeting ineffectiveness [3, 15, 32, 37, 38, 47].[5] The damage cascades across the entire meeting lifecycle [8, 15, 37, 38, 50]: preparation suffers, moderation lacks focus, discussions drift, and post-meeting effort increases. Research on AI-assisted prospective reflection offers a counterpoint. Prompting people at meeting creation to articulate purpose, success conditions, and challenges makes implicit goals explicit, clarifies priorities, and surfaces unknowns [37]. The effects persist: people share better agendas, request preparation from others, arrive with clearer intent, and conduct more focused meetings [37]. A small upstream investment in goal articulation produces disproportionate downstream value for everyone involved and every system that touches the meeting [47]. Once encoded, meeting goals ground downstream agents: a recap agent can assess what was achieved rather than merely what was said, a scheduling agent can reason about whether the right people are invited, and a task agent can link action items to the outcomes they serve [47, 50].

[2]https://asana.com/product/goals
[3]https://clickup.com/features/goals
[4]https://www.atlassian.com/software/jira/goals
[5]Goals are distinct from agendas, which describe *what you will do*, not *why*.

Meetings illustrate a general pattern. Explicit goals reduce the ambiguity that agents otherwise navigate through guesswork, improving the relevance and evaluability of their contributions. Yet almost no system deliberately collects goal information at the point where it can have the most value. This is both the opportunity and the design challenge.

5 Design Requirements for Goal-Oriented Systems

The idea that knowledge work benefits from explicit goal structures is not new. Group Support Systems [11], workflow management platforms, and successive waves of "rational" coordination tools have all tried to systematise the intentions behind collaborative work [36, 51]. Most foundered on the same problem: real goals are fuzzy, fluid, and context-dependent, while the systems designed to capture them demanded premature precision [4]. Indeed, the GOMS model in HCI assumes that users already have clear goals [22]. As Schmidt and Bannon argued, organisational models in CSCW systems should not be rigid specifications but "resources for competent and responsible workers" that users interpret, adapt, and override as needed [36]. Systems that imposed brittle goal structures on the "inexhaustible multiplicity of reality" were inevitably circumvented or abandoned.

Generative AI changes this equation. Large language models can work with goals expressed in natural language, at whatever level of specificity the user can manage, without requiring formal decomposition up front [49]. They can interpret ambiguous intent [23], ask clarifying questions [37], infer implicit goals from context [13], and track how goals evolve through conversation and action [8, 9]. For the first time, it is feasible to build goal support that respects the fuzziness Schmidt and Bannon [36] identified as inescapable while still providing the representational backbone for coordination, evaluation, and agent orchestration. Goals can now be accommodated computationally. The open question is how to treat goal representations as living resources, open to interpretation and revision, not fixed contracts. With this in mind, we identify six capabilities that a goal-oriented system should support:

(1) **Goal articulation.** Systems should help users and AI agents distinguish outcomes from outputs, decompose objectives into sub-goals, and reason across goal altitudes. This means balancing the precision needed to specify work with the ambiguity needed for reflection [4]. AI should play a substantive collaborative role: interrogating whether a goal is coherent, feasible, and aligned with what a team is actually doing, not just polishing phrasing [20].
(2) **Goals as first-class representations.** Goals should be structured, persistent objects that link to people, agents, artefacts, and workflows. They should accommodate fuzziness, evolution, and cross-level relationships. Making goals representationally first-class is what allows them to travel across tools and contexts.
(3) **Alignment and misalignment surfacing.** Systems should show users how their goals connect to others' and to higher organisational levels. They should detect implicit or emerging goals, surface misalignment proactively, and support structured negotiation to resolve conflicts.
(4) **Goal provenance and history.** Systems should track how goals evolve: who changed them, when, and why. This supports accountability and organisational learning, and gives AI agents richer context for assessing whether current work still serves its original purpose.
(5) **Outcome simulation.** Users should be able to anticipate how decisions cascade across goal dependencies: lightweight "what-if" projections that move goal management from retrospective tracking to prospective reasoning.
(6) **Situated goal support in workflows.** Goal support should be woven into workflows, not confined to dashboards: appearing as situated cues in the document being drafted, the meeting being run, or the code being reviewed.

This extends to agent orchestration. AI agents should receive delegated goals and report progress against them, not merely against tasks completed.

6 Conclusion

The ceiling on human-AI collaboration is not the capability of the AI; it is the clarity of the goals we give it. Raising that ceiling requires collaborative systems that help people articulate goals, encode them as shared representations, and thread them through workflows where humans and agents coordinate. The research agenda has three priorities. First, designing goal articulation interfaces that support reflective, iterative thinking beyond form-filling, and ensure mutual intelligibility for humans and AI. Second, building interoperability layers that let goal representations travel across tools and contexts. Third, developing goal-oriented experiences (and evaluation frameworks) that drive collaboration towards outcomes over outputs.

References

[1] Ruth V. Aguilera, Alfredo De Massis, Riccardo Fini, and Silvio Vismara. 2024. Organizational Goals, Outcomes, and the Assessment of Performance: Reconceptualizing Success in Management Studies. *Journal of Management Studies* 61, 1 (2024), 1–36. doi:10.1111/joms.12994 _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/joms.12994.

[2] Anthropic. 2025. Ralph Wiggum Plugin for Claude Code. https://github.com/anthropics/claude-code/tree/main/plugins/ralph-wiggum. Accessed: 2026-02-19.

[3] Henning Bang, Synne L. Fuglesang, Mariann R. Ovesen, and Dag Erik Eilertsen. 2010. Effectiveness in Top Management Group Meetings: The Role of Goal Clarity, Focused Communication, and Learning Behavior. *Scandinavian Journal of Psychology* 51, 3 (2010), 253–261. doi:10.1111/j.1467-9450.2009.00769.x

[4] Gagan Bansal, Jennifer Wortman Vaughan, Saleema Amershi, Eric Horvitz, Adam Fourney, Hussein Mozannar, Victor Dibia, and Daniel S. Weld. 2024. Challenges in Human-Agent Communication. doi:10.48550/arXiv.2412.10380 arXiv:2412.10380 [cs].

[5] Gagan Bansal, Tongshuang Wu, Joyce Zhou, Raymond Fok, Besmira Nushi, Ece Kamar, Marco Tulio Ribeiro, and Daniel S. Weld. 2021. Does the Whole Exceed Its Parts? The Effect of AI Explanations on Complementary Team Performance. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM. doi:10.1145/3411764.3445717

[6] Jenna L. Butler, Thomas Zimmermann, and Christian Bird. 2024. Objectives and Key Results in Software Teams: Challenges, Opportunities and Impact on Development. In *Proceedings of the 46th International Conference on Software Engineering: Software Engineering in Practice*. ACM, Lisbon Portugal, 358–368. doi:10.1145/3639477.3639747

[7] Mert Cemri, Melissa Z. Pan, Shuyi Yang, Lakshya A. Agrawal, Bhavya Chopra, Rishabh Tiwari, Kurt Keutzer, Aditya Parameswaran, Dan Klein, Kannan Ramchandran, Matei Zaharia, Joseph E. Gonzalez, and Ion Stoica. 2025. Why Do Multi-Agent LLM Systems Fail? doi:10.48550/arXiv.2503.13657 arXiv:2503.13657 [cs].

[8] Xinyue Chen, Lev Tankelevitch, Rishi Vanukuru, Ava Elizabeth Scott, Payod Panda, and Sean Rintel. 2025. Are We On Track? AI-Assisted Active and Passive Goal Reflection During Meetings. In *CHI Conference on Human Factors in Computing Systems (CHI '25)*. ACM, Yokohama, Japan, 22 pages. doi:10.1145/3706598.3714052

[9] Adam Coscia, Shunan Guo, Eunyee Koh, and Alex Endert. 2025. OnGoal: Tracking and Visualizing Conversational Goals in Multi-Turn Dialogue with Large Language Models. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology*. 1–18. doi:10.1145/3746059.3747746 arXiv:2508.21061 [cs].

[10] Allan Dafoe, Yoram Bachrach, Gillian Hadfield, Eric Horvitz, Kate Larson, and Thore Graepel. 2021. Cooperative AI: machines must learn to find common ground. *Nature* 593, 7857 (May 2021), 33–36. doi:10.1038/d41586-021-01170-0 Bandiera_abtest: a Cg_type: Comment Number: 7857 Subject_term: Machine learning, Computer science, Society, Technology, Sociology, Human behaviour.

[11] Gert-Jan de Vreede, Robert M. Davison, and Robert O. Briggs. 2003. How a silver bullet may lose its shine. *Commun. ACM* 46, 8 (Aug. 2003), 96–101. doi:10.1145/859670.859676

[12] John Doerr. 2018. *Measure What Matters: How Google, Bono, and the Gates Foundation Rock the World with OKRs*. Portfolio/Penguin, New York, NY, USA.

[13] Ian Drosos, Jack Williams, Advait Sarkar, and Nicholas Wilson. 2024. Dynamic Prompt Middleware: Contextual Prompt Refinement Controls for Comprehension Tasks. doi:10.48550/arXiv.2412.02357 arXiv:2412.02357 [cs].

[14] Peter F. Drucker. 1967. *The Effective Executive*. Harper & Row, New York, NY, USA.

[15] Jennifer L. Geimer, Desmond J. Leach, Justin A. DeSimone, Steven G. Rogelberg, and Peter B. Warr. 2015. Meetings at work: Perceived effectiveness and recommended improvements. *Journal of Business Research* 68, 9 (Sept. 2015), 2015–2026. doi:10.1016/j.jbusres.2015.02.015

[16] Peter M. Gollwitzer and Paschal Sheeran. 2025. Psychology of Planning. *Annual Review of Psychology* 76, 1 (Jan. 2025), 303–328. doi:10.1146/annurev-psych-021524-110536

[17] Francis Green. 2012. Employee Involvement, Technology and Evolution in Job Skills: A Task-Based Analysis. *Industrial and Labor Relations Review* 65, 1 (2012), 36–67. https://www.jstor.org/stable/41343664

[18] John Hattie and Helen Timperley. 2007. The Power of Feedback. *Review of Educational Research* 77, 1 (March 2007), 81–112. doi:10.3102/003465430298487

[19] Clara Herkenrath, Gerrit Hoeborn, and Volker Stich. 2023. Why Companies Fail With Objectives And Key Results: An Analysis Of Implementation Frameworks. (2023). doi:10.15488/13449 Publisher: Hannover : publish-Ing..

[20] Eric Horvitz. 1999. Principles of mixed-initiative user interfaces. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems (CHI '99)*. Association for Computing Machinery, New York, NY, USA, 159–166. doi:10.1145/302979.303030

[21] Marleen Huysman. 2025. Studying AI in the Wild: Reflections from the AI@Work Research Group. *Journal of Management Studies* n/a, n/a (2025). doi:10.1111/joms.70021 _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/joms.70021.

[22] Bonnie E John and David E Kieras. 1996. Using GOMS for user interface design and evaluation: Which technique? *ACM Transactions on Computer-Human Interaction (TOCHI)* 3, 4 (1996), 287–319.

[23] Yoonsu Kim, Brandon Chin, Kihoon Son, Seoyoung Kim, and Juho Kim. 2025. IntentFlow: Interactive Support for Communicating Intent with LLMs in Writing Tasks. doi:10.48550/arXiv.2507.22134 arXiv:2507.22134 [cs].

[24] Gary P. Latham and Edwin A. Locke. 1991. Self-Regulation Through Goal Setting. *Organizational Behavior and Human Decision Processes* 50, 2 (1991), 212–247. doi:10.1016/0749-5978(91)90021-K

[25] Edwin Locke and Gary Latham. 2005. Goal-Setting Theory. In *Organizational Behavior 1*. Routledge. Num Pages: 25.

[26] Edwin A. Locke and Gary P. Latham. 2002. Building a Practically Useful Theory of Goal Setting and Task Motivation: A 35-Year Odyssey. *American Psychologist* 57, 9 (2002), 705–717. doi:10.1037/0003-066X.57.9.705

[27] André N. Meyer, Gail C. Murphy, Thomas Zimmermann, and Thomas Fritz. 2021. Enabling Good Work Habits in Software Developers through Reflective Goal-Setting. *IEEE Transactions on Software Engineering* 47, 9 (Sept. 2021), 1872–1885. doi:10.1109/TSE.2019.2938525 Conference Name: IEEE Transactions on Software Engineering.

[28] Microsoft. 2023. Work Trend Index: Will AI Fix Work? https://www.microsoft.com/en-us/worklab/work-trend-index/will-ai-fix-work Accessed: 2026-02-11.

[29] Microsoft. 2025. *Microsoft 2025 Work Trend Index Annual Report*. Technical Report. https://www.microsoft.com/en-us/worklab/work-trend-index/2025-the-year-the-frontier-firm-is-born

[30] Aruna Ranganathan and Xingqi Maggie Ye. 2026. AI Doesn't Reduce Work—It Intensifies It. *Harvard Business Review* (Feb. 2026). https://hbr.org/2026/02/ai-doesnt-reduce-work-it-intensifies-it Section: AI and machine learning.

[31] Matthew K. Robison, Nash Unsworth, and Gene A. Brewer. 2021. Examining the effects of goal-setting, feedback, and incentives on sustained attention. *Journal of Experimental Psychology: Human Perception and Performance* 47, 6 (June 2021), 869–891. doi:10.1037/xhp0000926

[32] N.C. Romano and J.F. Nunamaker. 2001. Meeting analysis: findings from research and practice. In *Proceedings of the 34th Annual Hawaii International Conference on System Sciences*. 13 pp.–. doi:10.1109/HICSS.2001.926253

[33] Advait Sarkar. 2023. Exploring Perspectives on the Impact of Artificial Intelligence on the Creativity of Knowledge Work: Beyond Mechanised Plagiarism and Stochastic Parrots. In *Proceedings of the 2nd Annual Meeting of the Symposium on Human-Computer Interaction for Work*. ACM, Oldenburg Germany, 1–17. doi:10.1145/3596671.3597650

[34] Suproteem K. Sarkar. 2025. AI Agents, Productivity, and Higher-Order Thinking: Early Evidence From Software Development. doi:10.2139/ssrn.5713646

[35] Sam Schillace. 2026. Attention and collaboration in the AI world. https://sundaylettersfromsam.substack.com/p/attention-and-collaboration-in-the?utm_medium=email&action=share

[36] Kjeld Schmidt and Liam Bannon. 1992. Taking CSCW Seriously: Supporting Articulation Work. *Computer Supported Cooperative Work (CSCW)* 1, 1–2 (1992), 7–40. doi:10.1007/BF00752449

[37] Ava Elizabeth Scott, Lev Tankelevitch, Payod Panda, Rishi Vanukuru, Xinyue Chen, and Sean Rintel. 2025. What Does Success Look Like? Catalyzing Meeting Intentionality with AI-Assisted Prospective Reflection. In *CHIWORK '25: Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work*. ACM, Amsterdam, Netherlands, 33 pages. doi:10.1145/3729176.3729204

[38] Ava Elizabeth Scott, Lev Tankelevitch, and Sean Rintel. 2024. Mental Models of Meeting Goals: Supporting Intentionality in Meeting Technologies. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. ACM, Honolulu, HI, USA, 17 pages. doi:10.1145/3613904.3642670

[39] Alexandra Spitz-Oener. 2006. Technical Change, Job Tasks, and Rising Educational Demands: Looking outside the Wage Structure. *Journal of Labor Economics* 24, 2 (2006), 235–270. https://ideas.repec.org//a/ucp/jlabec/v24y2006i2p235-270.html

[40] Philipp Spitzer, Joshua Holstein, Patrick Hemmer, Michael Vössing, Niklas Kühl, Dominik Martin, and Gerhard Satzger. 2025. Human Delegation Behavior in Human-AI Collaboration: The Effect of Contextual Information. *Proc. ACM Hum.-Comput. Interact.* 9, 2 (May 2025), CSCW101:1–CSCW101:28. doi:10.1145/3710999

[41] Anselm Strauss. 1985. Work and the Division of Labor. *The Sociological Quarterly* 26, 1 (1985), 1–19. doi:10.1111/j.1533-8525.1985.tb00212.x _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1533-8525.1985.tb00212.x.

[42] Viktoria Stray, Jan Henrik Gundelsby, Rasmus Ulfsnes, and Nils Brede Moe. 2022. How agile teams make Objectives and Key Results (OKRs) work. In *Proceedings of the International Conference on Software and System Processes and International Conference on Global Software Engineering*. ACM, Pittsburgh PA USA, 104–109. doi:10.1145/3529320.3529332

[43] Hari Subramonyam, Divy Thakkar, Andrew Ku, Juergen Dieber, and Anoop K. Sinha. 2025. Prototyping with Prompts: Emerging Approaches and Challenges in Generative AI Design for Collaborative Software Teams. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–22. doi:10.1145/3706598.3713166

[44] Lucille Alice Suchman. 1987. *Plans and situated actions: The problem of human-machine communication.* Cambridge university press.

[45] Lev Tankelevitch, Elena L. Glassman, Jessica He, Aniket Kittur, Mina Lee, Srishti Palani, Advait Sarkar, Gonzalo Ramos, Yvonne Rogers, and Hari Subramonyam. 2025. Understanding, Protecting, and Augmenting Human Cognition with Generative AI: A Synthesis of the CHI 2025 Tools for Thought Workshop. doi:10.48550/arXiv.2508.21036 arXiv:2508.21036 [cs].

[46] Lev Tankelevitch, Viktor Kewenig, Auste Simkute, Ava Elizabeth Scott, Advait Sarkar, Abigail Sellen, and Sean Rintel. 2024. The Metacognitive Demands and Opportunities of Generative AI. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–24. doi:10.1145/3613904.3642902

[47] Lev Tankelevitch, Ava Elizabeth Scott, Nagaravind Challakere, Payod Panda, and Sean Rintel. 2026. Nudging Attention to Workplace Meeting Goals: A Large-Scale, Preregistered Field Experiment. doi:10.1145/3772318.3791199 arXiv:2602.16939 [cs].

[48] Nenad Tomašev, Matija Franklin, and Simon Osindero. 2026. Intelligent AI Delegation. doi:10.48550/arXiv.2602.11865 arXiv:2602.11865 [cs].

[49] Priyan Vaithilingam, Munyeong Kim, Frida-Cecilia Acosta-Parenteau, Daniel Lee, Amine Mhedhbi, Elena L. Glassman, and Ian Arawjo. 2025. Semantic Commit: Helping Users Update Intent Specifications for AI Memory at Scale. doi:10.48550/arXiv.2504.09283 arXiv:2504.09283 [cs].

[50] Rishi Vanukuru, Payod Panda, Xinyue Chen, Ava Elizabeth Scott, Lev Tankelevitch, and Sean Rintel. 2025. Designing Interfaces that Support Temporal Work Across Meetings with Generative AI. In *Designing Interactive Systems Conference (DIS '25)*. ACM, Funchal, Portugal, 21 pages. doi:10.1145/3715336.3735833

[51] Amy Voida, Ellie Harmon, and Ban Al-Ani. 2011. Homebrew databases: complexities of everyday information management in nonprofit organizations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. Association for Computing Machinery, New York, NY, USA, 915–924. doi:10.1145/1978942.1979078

[52] Monika Westphal, Patrick Hemmer, Michael Vössing, Max Schemmer, Sebastian Vetter, and Gerhard Satzger. 2025. Towards Understanding AI Delegation: The Role of Self-Efficacy and Visual Processing Ability. *ACM Trans. Interact. Intell. Syst.* 15, 1 (Feb. 2025), 5:1–5:24. doi:10.1145/3696423

[53] Christina Wodtke. 2016. *Radical Focus.* Cucina Media.