

WORKING PAPER 196

CONTEXT-MAXXING

A PATH TO COGNITIVE AGENCY WITH GENERATIVE AI

JACOB TAYLOR
KERSHLIN KRISHNA

Context-maxxing: A path to cognitive agency with generative AI

Jacob Taylor and Kershlin Krishna

APRIL 2026
Working Paper

About the Center for Sustainable Development

The Center for Sustainable Development (CSD) generates cutting-edge research and empowers collaborative action to tackle the world’s interwoven economic, social, and environmental challenges in all societies and at all scales. To learn more visit: www.brookings.edu/SustainableDevelopment

About the Global Economy and Development

Founded in 2006, the Global Economy and Development program at the Brookings Institution aims to play its part to ensure that the future of globalization is one of inclusive growth and shared prosperity. With an interdisciplinary team of experts, Global provides thought-leadership, cutting-edge research, and innovative policy solutions to achieve a more equitable global economic system for sustainable prosperity, drawing on the core strengths of Brookings—authoritativeness, independence, depth of practical expertise, and unparalleled convening power. For more, visit www.brookings.edu/Global.

Acknowledgment

The authors thank Homi Kharas and Caren Grown for review and editorial oversight, as well as Scott Page, Christoph Riedl, and Elham Tabassi, who thoughtfully reviewed earlier versions of this paper. The authors also want to acknowledge a small but growing community of lead users, whose (largely) open experimentation with open-source agent harness approaches helped inform the playbook. With thanks also to Laura Martínez Rodríguez, Clea McElwain, and Kendall Trelegan for their evaluation work and contributions to the Brookings application. Finally, thanks to Liv Schaffer for fact-checking and Junjie Ren for production and communications support.

Disclosure

This paper was developed in accordance with the “context-maxxing” approach described herein. Generative AI models (provided by Anthropic, Gemini, OpenAI, and PublicAI) hosted in an open-source harness (OpenClaw) assisted in synthesizing expert analysis and drafting initial outputs. All final content was revised by the authors and reviewed for factual accuracy prior to publication.

Support for this publication was generously provided by The Rockefeller Foundation and the Sloan Foundation. Brookings is committed to quality, independence, and impact. Brookings is supported by a diverse array of funders. In line with Brookings values and policies, each Brookings publication represents the sole views of its author(s).

Executive summary

In January 2026, [OpenClaw](#)—open-source software that allows users to interact with AI models within their own computing environments—went viral, surpassing 100,000 GitHub stars within weeks. This bottom-up movement introduced a fundamentally different deployment paradigm for generative AI and a glimpse of what it could mean to center people’s interests within it. Using OpenClaw and other open-source agent harnesses, millions of users globally are now exercising greater control over the information that they bring to their interactions with AI—domain knowledge, intentionality, and judgment—referred to in the industry as “context.”

This paper codifies common elements of emerging user-controlled, open-source AI deployments “context-maxxing.”¹ The claim, motivated by analysis of the cognitive effects of existing proprietary AI deployment infrastructures and the incentives that drive their design, is that context-maxxing could increase cognitive agency (a capacity for people to think and act with generative AI in ways that support control, efficacy, and mastery) relative to default vendor-controlled interfaces like ChatGPT or Claude.ai.

This working paper is intended for policymakers and decisionmakers interested in how generative AI can support new forms of value creation and shared problem-solving. It is structured in three parts. Part 1 analyzes the relationship between existing generative AI deployment paradigms and cognitive agency as a basis for proposing context-maxxing as a promising alternative approach. Part 2 presents a playbook for context-maxxing designed to inform decisionmakers of the core digital infrastructure building blocks and emerging competencies associated with user-controlled open-source AI deployments, and how they could support cognitive agency. Part 3 discusses policy implications and future directions for empirical research.

This work motivates a research agenda to test whether the way generative AI is deployed shapes human cognitive agency. The paper’s analysis suggests that computing environments that prioritize user control over context could lay the foundation for entirely new human expertise and capabilities for shared problem-solving with AI. This assertion in turn raises questions about how public policy interventions can lower barriers to context-maxxing for people everywhere—including in sustainable development contexts where the infrastructure and resources to do so remain unevenly distributed.

¹ The “-maxxing” suffix—internet shorthand for obsessive optimization—is used deliberately in an attempt to reclaim the meme to support human interests. In live AI industry debates, “token-maxxing” urges firms to maximize employee AI consumption; “outcome-maxxing” counters that firms should optimize for what AI produces. Context-maxxing enters this debate by asking who controls the context through which those outputs are produced. The term is also a gesture toward more agile and culturally legible modes of policy research, prototyped at the edge of the systems where generative AI is being deployed.

Part 1. A path to human cognitive agency with generative AI

Generative artificial intelligence (AI) is transforming human activity at an unprecedentedly intimate layer of human cognition. Over the last 15 years, increasingly powerful machine learning algorithms integrated into internet platform applications have reorganized human behavior at a societal scale—reducing the cost of communication, computation, and commerce, while at the same time sorting and manipulating attention,² reducing competition in markets,³ and expanding and compromising information ecosystems. Even before generative AI, human activity was already unfolding in an algorithmic age, in which technology shapes human behavior and information ecosystems.⁴

Deployed on top of this infrastructure, generative AI is now beginning to mediate thought itself. Large language models (LLMs) generate probabilistic outputs through an interaction between user-provided information—originally referred to as prompts, now more broadly termed “context”—and models trained on internet-scale volumes of human-generated text, audio, and video.^{5,6} Unlike earlier AI systems, LLMs are available through free or low-cost consumer products, operable through ordinary language, and applicable across virtually any domain. AI systems now engage meaningfully and in real time with users’ reasoning, judgment, and intentionality—capacities long considered intrinsically human.

Now that humans are thinking and acting in partnership with AI—in roles that may include coworker, coach, augments, critic, and, in some cases, autonomous agent—a key challenge for public policy is understanding whether or not this partnership serves human interests. This means first defining human interests at the site of activity at which this new partnership is taking shape—human cognition itself. Second, it requires analyzing the relationship between generative AI’s deployment and human interests. Third, it necessitates identifying practical and timely policy responses to ensure that these relationships serve human interests. This paper addresses each step in turn.

² Chen, Yi, and Li, Fei, and Preuss, Marcel, “Algorithmic Attention and Content Creation on Social Media Platforms” (March 17, 2025), available at SSRN: <https://ssrn.com/abstract=5182754>.

³ Lynn, Barry C., “The Big Tech Extortion Racket: How Google, Amazon, and Facebook Control Our Lives,” *Harper’s Magazine*, September 2020.

⁴ John Danaher et al., “Algorithmic governance: Developing a research agenda through the power of collective intelligence,” *Big Data & Society* 4, no. 2 (2017), <https://doi.org/10.1177/2053951717726554>

⁵ “Context Engineering,” Gartner, accessed April 15, 2026, <https://www.gartner.com/en/articles/context-engineering>.

⁶ Banh, L., and Strobel, G., “Generative Artificial Intelligence,” *Electronic Markets* 33, 63 (2023), <https://doi.org/10.1007/s12525-023-00680-1>.

Cognitive agency in human-AI collaboration

Recent attempts to elevate human cognitive interests and freedoms amid the rapid global diffusion of generative AI include proposals for cognitive sovereignty⁷ (the autonomy, reflectiveness, and resilience people need to maintain control over their reasoning) and cognitive liberty⁸ (rights-based protections against interference with mental self-determination). While fully affirming these principles of ensuring human interests and freedoms from technological harms, here the aim is to extend these discussions by emphasizing opportunities for people and collectives to think and act in a dynamic partnership with AI systems in ways that expand human potential and increase capabilities for shared problem-solving.⁹

This paper, therefore, proposes cognitive agency as the capacity for people to think and act with AI in ways that support their control, efficacy, and mastery.

This account is grounded in existing work that identifies agency as foundational to human and sustainable development outcomes: from Amartya Sen's account of human development as the expansion of people's freedom to live lives they have reason to value,^{10,11} through Naila Kabeer's concept of empowerment as resources, agency, and achievements,¹² and Elinor Ostrom's demonstration that durable collective action over shared resources depends on the agency of the communities that govern them.¹³ Building on Floridi's work on informational agency, cognitive agency further emphasizes the ability of people (encompassing individuals and collectives) to think and act within a digital environment shared with AI systems.¹⁴

⁷ Mahadi Mokbul Ali, "Cognitive Sovereignty: A Theory and Initial Validation of Human Autonomy in Algorithmic Decision Systems," Research Square, version 1, posted March 17, 2026, <https://www.researchsquare.com/article/rs-9145237/v1>.

⁸ Courtney C. Radsch, "The Battle for Cognitive Liberty in the Age of Corporate AI," Tech Policy Press, January 6, 2026, <https://www.techpolicy.press/the-battle-for-cognitive-liberty-in-the-age-of-corporate-ai/>.

⁹ United Nations Development Programme. Human Development Report 2025: A Matter of Choice – People and Possibilities in the Age of AI. New York: UNDP, 2025. <https://hdr.undp.org/content/human-development-report-2025>.

¹⁰ Amartya Sen, *Commodities and Capabilities* (Amsterdam: North-Holland, 1985); Amartya Sen, *Development as Freedom* (New York: Alfred A. Knopf, 1999).

¹¹ Pelenc, Jérôme, and Catherine Ballet. "Is Amartya Sen's Sustainable Freedom a Broader Vision of Sustainability?" *Ecological Economics* 105 (September 2014): 1–7. <https://doi.org/10.1016/j.ecolecon.2014.05.010>

¹² Kabeer, Naila. "Resources, Agency, Achievements: Reflections on the Measurement of Women's Empowerment." *Development and Change* 30, no. 3 (1999): 435–464. <https://doi.org/10.1111/1467-7660.00125>.

¹³ Ostrom, Elinor. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press, 1990.

¹⁴ Luciano Floridi, *The Ethics of Information* (Oxford: Oxford University Press, 2013).

Drawing these literatures together with behavioral and socio-cognitive research on preconditions for individual and collective agency, three analytic dimensions of cognitive agency with generative AI can be assessed:

- **Control:** the capacity to shape the informational environment in which thinking and acting with AI takes place. This includes what information is captured and stored, how it is represented, who has access to it, and on what terms. This dimension reflects social psychology research that suggests individual agency is associated with predictability, responsibility, and accountability over¹⁵—as well as metacognitive monitoring of¹⁶—one’s own thinking or action, while collective agency is associated with equal participation and mutual information flow.¹⁷ Additionally, control in a digital environment also means technical capabilities of portability or interoperability—whether what a person or collective builds remains with them when they move between tools, platforms, or institutions.¹⁸
- **Efficacy:** the capacity to think and act effectively with AI to pursue goals and address shared challenges. For individuals, this includes perceptions of self-efficacy built through immediate, attributable feedback.^{19,20} For collectives, it is the shared efficacy required for problem-solving across diverse perspectives, including the need for cognitive diversity,²¹ social perceptiveness (or Theory of Mind),²² and shared understanding of goals.
- **Mastery:** the capacity to experience, learn, and grow durable capability through human-AI interaction over time. At the individual level, mastery is associated with optimal performance states and learning effects.^{23,24} At the collective level, it

¹⁵Deci, Edward L., and Richard M. Ryan. “The ‘What’ and ‘Why’ of Goal Pursuits: Human Needs and the Self-Determination of Behavior.” *Psychological Inquiry* 11, no. 4 (2000): 227–68.

¹⁶ Flavell, John H. “Metacognition and Cognitive Monitoring: A New Area of Cognitive-Developmental Inquiry.” *American Psychologist* 34, no. 10 (1979): 906–11.

¹⁷ Woolley, Anita Williams, Christopher F. Chabris, Alex Pentland, Nada Hashmi, and Thomas W. Malone. “Evidence for a Collective Intelligence Factor in the Performance of Human Groups.” *Science* 330, no. 6004 (2010): 686–88.

¹⁸ Eaves, David. “The Path to a Sovereign Tech Stack Is via a Commodified Tech Stack.” *Tech Policy Press*, December 15, 2025. <https://techpolicy.press/the-path-to-a-sovereign-tech-stack-is-via-a-commodified-tech-stack/>.

¹⁹ Bandura, Albert. “Social Cognitive Theory: An Agentic Perspective.” *Annual Review of Psychology* 52 (2001): 1-26.

²⁰ Csikszentmihalyi, Mihaly. *Flow: The Psychology of Optimal Experience*. New York: Harper & Row, 1990.

²¹ Page, Scott E. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton, NJ: Princeton University Press, 2007.

²² Christoph Riedl, Young Ji Kim, Pranav Gupta, Thomas W. Malone, and Anita Williams Woolley, “Quantifying Collective Intelligence in Human Groups,” *Proceedings of the National Academy of Sciences of the United States of America* 118, no. 21 (May 25, 2021): e2005737118, <https://doi.org/10.1073/pnas.2005737118>.

²³ Ericsson, K. Anders, Ralf Th. Krampe, and Clemens Tesch-Römer. “The Role of Deliberate Practice in the Acquisition of Expert Performance.” *Psychological Review* 100, no. 3 (1993): 363–406.

²⁴ Csikszentmihalyi, Mihaly. *Flow: The Psychology of Optimal Experience*. New York: Harper & Row, 1990.

tracks what cognitive anthropologists call cumulative culture: the durable shared infrastructure through which capability passes from one cohort of practitioners to the next.^{25,26}

The “informational environment” in which human thinking and acting with AI takes place is shaped by AI’s deployment infrastructure, or the sociotechnical stack of hardware, software, data pipelines, interfaces, and governance arrangements that make generative AI available to people in concrete settings.²⁷ As explained below, AI deployment infrastructure is increasingly shaped by incentives to control context—or the information that users provide to generative AI systems.

Context is the new data

Because LLMs are non-deterministic, stateless, and do not inherently accumulate knowledge over time, their performance depends critically on the richness and accuracy of information provided in each exchange.^{28,29} As generative AI increasingly supports judgment-dependent domains where outputs cannot always be mechanically verified (e.g., knowledge work), its usefulness depends less on model capability alone and more on a partnership between human-provided information—or what is now referred to in the AI industry as context—and the inference power of generative AI systems.^{30,31}

The explosion of cloud storage and machine-learning infrastructure underpinning today’s internet platforms means that people now disclose—intentionally or otherwise—orders of magnitude more of their context into digital environments than was previously imaginable. Generative AI deployments are accelerating this trend on a radically compressed timeline.

²⁵ Henrich, Joseph. *The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*. Princeton, NJ: Princeton University Press, 2016.

²⁶ Muthukrishna, Michael, and Joseph Henrich. “Innovation in the Collective Brain.” *Philosophical Transactions of the Royal Society B: Biological Sciences* 371, no. 1690 (2016): 20150192.

²⁷ Maas, Matthijs M. “Sociotechnical Change: AI as Regulatory Rationale and Target.” In *Architectures of Global AI Governance: From Technological Change to Human Choice*, edited by Wulf A. Kaal and Francesco Paolo Patti, chapter 6. Oxford: Oxford University Press, 2025.
<https://doi.org/10.1093/9780191988455.003.0006>.

²⁸ Vaswani, Ashish, et al., “Attention Is All You Need,” arXiv preprint, first posted June 12, 2017, <https://arxiv.org/abs/1706.03762>.

²⁹ Sahoo, Pranab, et al., “A Systematic Survey of Prompt Engineering in Large Language Models: Techniques and Applications,” arXiv preprint, 2024, <https://arxiv.org/abs/2402.07927>.

³⁰ Guszczka, James and Danks, David and Fox, Craig R. and Hammond, Kristian J. and Ho, Daniel E. and Imas, Alex and Landay, James and Levi, Margaret and Logg, Jennifer and Picard, Rosalind W. and Raghavan, Manish and Stanger, Allison and Ugolnik, Zachary and Woolley, Anita Williams, *Hybrid Intelligence: A Paradigm for More Responsible Practice* (October 12, 2022). Available at SSRN: <https://ssrn.com/abstract=4301478> or <http://dx.doi.org/10.2139/ssrn.4301478>

³¹ Christoph Riedl, Saiph Savage, and Josie Zvelebilova, “Cognitive Spillover in Human-AI Teams,” *ACM Transactions on Computer-Human Interaction*, Just Accepted (April 2026), <https://doi.org/10.1145/3805039>.

Systems have evolved from operating with very narrow context windows and sparse prompts to handling far more information via longer context windows, more sophisticated retrieval-augmented generation, and tool calls (the ability for AI to perform functions through external applications).³²

As a result, user interactions with AI have evolved from a paradigm of prompt engineering into context engineering: an ongoing process of supplying increasingly richer, more situated information—domain knowledge, reasoning patterns, constraints, and intent—that increasingly autonomous or agentic AI systems synthesize and act upon (see Figure 1). This shift has driven step-changes in performance and capability, owing to an ability to weave high volumes of information (prompts, documents, code, prior interactions, and external tools like web search) into a single reasoning trace.³³

For model builders, too, context is a strategic asset. Subject to users' terms of service, user context is captured and used to improve foundation models (and the surrounding software that enables their deployment) through base model training, fine-tuning, and reinforcement learning.^{34,35,36} The result is a context flywheel: Successful deployment generates context, context improves the product, improved products attract more users, and more users generate richer context.³⁷ Recent shifts in venture-capital vocabulary from "data is the moat" to "taste is the moat" signal that AI firms and investors

³² Bergmann, Dave. "What Is a Context Window?" IBM. Accessed April 20, 2026.

<https://www.ibm.com/think/topics/context-window>

³³ Håkon Lønningdal Joren et al., "Sufficient Context: A New Lens on Retrieval-Augmented Generation Systems," in Proceedings of the 13th International Conference on Learning Representations (ICLR 2025) (Vienna, 2025), <https://arxiv.org/abs/2411.06037>.

³⁴ OpenAI. "How Your Data Is Used to Improve Model Performance." OpenAI, updated April 28, 2025. <https://openai.com/policies/how-your-data-is-used-to-improve-model-performance/>.

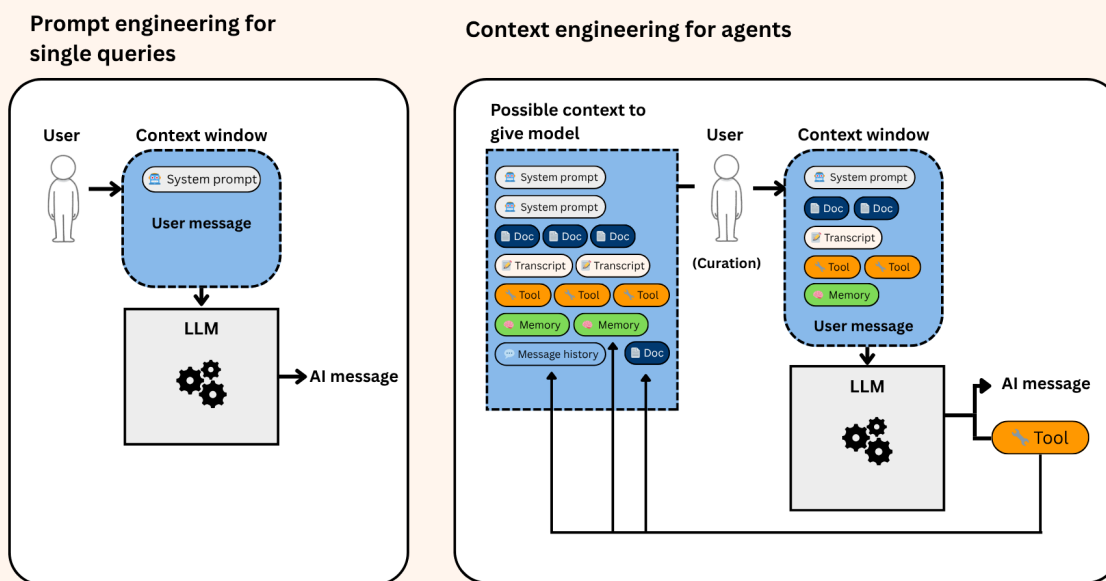
³⁵ The strength and directness of this loop varies substantially across firms, products, and user tiers; the clearest and most empirically documented cases to date involve free-tier consumer products, where user interactions are most frequently available for training by default. For example, [as of early 2026](#), free-tier users of OpenAI's ChatGPT constitute approximately 94% of the product's approximately 900 million weekly active users, and are subject to default opt-in to training use unless they manually adjust settings; comparable patterns hold for Google's Gemini free tier. Paid API and enterprise customers are typically, though not universally, excluded from training by contract.

³⁶ Indeed, the most striking recent performance gains in frontier language models in the past 18 months have come not from scaling foundation model training parameters but from [test-time scaling](#) (commonly referred to as inference), in which AI systems generate extended chains of internal deliberation before producing a response. This capability was itself learned through reinforcement learning over large collections of human-curated reasoning traces. In this specific and somewhat narrow sense, the "thinking" that reasoning models now perform internally is accumulated human cognitive context distilled into model weights and redeployed as autonomous capability.

³⁷ Patric Dubois. "The Context Flywheel: Why the Best AI Coding Teams Will Win on Context." TESSL Blog, accessed April 15, 2026. <https://tessl.io/blog/the-context-flywheel-why-the-best-ai-coding-teams-will-win-on-context/>

increasingly recognize context and the deployment infrastructure for controlling it as the primary vehicle for economic value capture through AI.^{38,39}

Figure 1. Difference between prompt engineering vs context engineering



Source: Adapted from [Anthropic](#).

The context that makes generative AI most productive for its users—their own domain knowledge, reasoning patterns, intentionality, and judgment—is the same resource that improves the models and deployment infrastructure they interact with. This suggests a potential synergy between user and vendor: Users making their context legible to generative AI systems should enrich both sides simultaneously because context, once codified, can be used by both parties without diminishing either.^{40,41}

Given how quickly these value capture dynamics are evolving, the relationship between the design and deployment of generative AI systems and cognitive agency deserves to be a more salient and urgent topic of public scrutiny and debate. As reviewed below, it

³⁸ Kyle Chayka, "Why Tech Bros Are Now Obsessed with Taste," *The New Yorker*, March 18, 2026, <https://www.newyorker.com/culture/infinite-scroll/why-tech-bros-are-now-obsessed-with-taste>.

³⁹ Josipa Majic, "VCs Say Context Graphs Might Be the Next Big Thing in AI," *Forbes*, April 3, 2026, <https://www.forbes.com/sites/josipamajic/2026/04/03/vcs-say-context-graphs-might-be-the-next-big-thing-in-ai/>

⁴⁰ Charles I. Jones and Christopher Tonetti, "Nonrivalry and the Economics of Data," NBER Working Paper No. 26260, September 2019, revised April 2020, <https://www.nber.org/papers/w26260>.

⁴¹ Anand V. Shah et al., "Robust AI Personalization Controls: The Human Context Protocol," SSRN Scholarly Paper 5403981 (September 7, 2025), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5403981.

appears that generative AI use is altering the cognitive processes through which people form judgments and develop expertise.^{42,43} The relationship between generative AI use and cognitive agency appears to depend on the extent to which users can control the deployment environment of these systems.

Generative AI use and cognitive agency

Most studies to date have overwhelmingly examined users interacting with AI through default proprietary interfaces (e.g., ChatGPT, Claude.ai, Gemini, or an equivalent chat-interface equivalent developed for experimental purposes). These interfaces represent a dominant mode of generative deployment in which users provide single-session inputs to an opaque system with limited control over how that context is stored, reused, or subsequently fed back into model improvement.

Under these default conditions of limited user control, LLM assistance has been shown to negatively impact perceived and actual efficacy and mastery. Studies show reduced neural engagement during composition tasks,⁴⁴ weakened memory consolidation and knowledge retention,⁴⁵ reduced persistence, weaker subsequent independent performance,⁴⁶ and diminishing reviewers' sense of authorship over outputs.⁴⁷ While cognitive offloading with AI could be adaptive (by outsourcing effort to AI tools to conserve energy for more meaningful tasks), studies have shown that the same tools lead to cognitive overloading (erosion of introspection and critical thinking can AI reduces critical thinking.^{48,49} These effects may have potential long-term implications for workers'

⁴² Anjali Singh, Karan Taneja, Zhitong Guan, and Avijit Ghosh, "Protecting Human Cognition in the Age of AI," arXiv preprint, last revised September 29, 2025, <https://arxiv.org/abs/2502.12447>.

⁴³ Christoph Riedl, Saiph Savage, and Josie Zvelebilova. 2026. Cognitive Spillover in Human-AI Teams. ACM Trans. Computer-Human Interaction. (April 2026). <https://doi.org/10.1145/3805039>.

⁴⁴ Nataliya Kosmyna et al., "Your Brain on ChatGPT: Accumulation of Cognitive Debt When Using an AI Assistant for Essay Writing Task," MIT Media Lab, June 7, 2025, <https://arxiv.org/pdf/2506.08872.pdf>.

⁴⁵ André Barcaui, ChatGPT as a cognitive crutch: Evidence from a randomized controlled trial on knowledge retention, *Social Sciences & Humanities Open*, Volume 12, 2025, 102287, <https://www.sciencedirect.com/science/article/pii/S2590291125010186>.

⁴⁶ Grace Liu, Brian Christian, Tsvetomira Dumbalska, Michiel A. Bakker, and Rachit Dubey, "AI Assistance Reduces Persistence and Hurts Independent Performance," arXiv preprint arXiv:2604.04721, revised April 7, 2026, doi:10.48550/arXiv.2604.04721.

⁴⁷ See, for example, Qin, Ying, Robert W. Smith, Mathieu Stillman, and Daniel M. Romero. "Generative AI Enhances Individual Creativity but Reduces the Collective Diversity of Novel Content." *Science Advances* 10, no. 28 (July 12, 2024): eadn5290. <https://pmc.ncbi.nlm.nih.gov/articles/PMC11244532/>.

⁴⁸ Georgios P. Georgiou, "ChatGPT Produces More 'Lazy' Thinkers: Evidence of Cognitive Engagement Decline," arXiv preprint arXiv:2507.00181, June 30, 2025, doi:10.48550/arXiv.2507.00181.

⁴⁹ Sandra Grinschgl, Frank Papenmeier, and Hauke Meyerhoff. (2021). Consequences of cognitive offloading: boosting performance but diminishing memory. *Q. J. Exp. Psychol.* 74, 1477–1496. doi: 10.1177/17470218211008060.

cognitive resilience, emotional regulation, and adaptive coping capacities.⁵⁰ The effects of cognitive erosion through generative AI use are already visible in the proliferation of AI workslop—where people losing both the interest and ability to verify the quality of AI-generated outputs is undermining authentic communication and economic productivity within organizations and, potentially, entire societies.^{51,52,53}

Research suggests that collectives of any scale relying on the same small set of LLM systems for sensemaking, analysis, and knowledge creation could drive what has been termed cognitive monoculture—a homogenization of patterns of thought and expression that could compromise cognitive diversity and resilience that underpin collective intelligence.⁵⁴ With LLMs acting as centralized knowledge systems controlled by a handful of firms, on the one hand, and users increasingly unable to exercise independent judgment due to cognitive offloading and overloading on the other, generative AI not only threatens to distort economic value creation, but it also introduces vulnerabilities to the reproduction of informational ecosystems that underpin democratic decisionmaking.^{55,56}

However, a growing parallel body of work suggests that generative AI use is associated with increases in cognitive agency and collective intelligence under conditions in which users exercise and experience greater control over their interactions with it. Most notably, structured and interactive prompting that elicits user reflection, planning, evaluation, and reasoning—compared to one-shot “just ask ChatGPT” queries—reverses the cognitive-offloading pattern and yields gains in critical reasoning.⁵⁷

Generative AI’s positive effects on cognitive performance thus appear to depend on interaction designs in which humans intentionally control both what they are doing with the system and how the system participates.⁵⁸ Lower perceptions of agency during

⁵⁰ Ginto Chirayath, K. Premamalini, and Jeena Joseph. “Cognitive Offloading or Cognitive Overload? How AI Alters the Mental Architecture of Coping.” *Frontiers in Psychology* 16 (2025): 1699320. <https://www.frontiersin.org/journals/psychology/articles/10.3389/fpsyg.2025.1699320/full>.

⁵¹ BetterUp Labs and Stanford Social Media Lab, “Workslop: The Hidden Cost of AI-Generated Busywork,” BetterUp Labs, September 20, 2025, <https://www.betterup.com/workslop>.

⁵² Christian Catalini, Xiang Hui, and Jane Wu, “Some Simple Economics of AGI,” arXiv preprint, February 25, 2026, <https://arxiv.org/abs/2602.20946>.

⁵³ Enrique Ide and Eduard Talamas, “Automation, Verification, and the Hollow Economy,” working paper, 2025

⁵⁴ Sourati, Zhivar, Ali S. Ziabari, and Mohammad Dehghani. “The Homogenizing Effect of Large Language Models on Human Expression and Thought.” *Trends in Cognitive Sciences* (advance online publication, 2026). <https://doi.org/10.1016/j.tics.2026.01.003>

⁵⁵ George, Rachel, and Ian Klaus. *AI and Democracy: Mapping the Intersections*. Washington, DC: Carnegie Endowment for International Peace, 2026. <https://carnegieendowment.org/research/2026/01/ai-and-democracy-mapping-the-intersections>

⁵⁶ Acemoglu, Daron, Tianyi Lin, Asuman Ozdaglar, and James Siderius. *How AI Aggregation Affects Knowledge*. NBER Working Paper 35036, April 2026. <https://www.nber.org/papers/w35036>

⁵⁷ Adam G. Fetterman, John B. Wilkerson, and Benjamin J. Blankenship, “From Offloading to Engagement: An Experimental Study on Structured Prompting and Critical Reasoning with Generative AI,” *Data* 10, no. 11 (2025): 172.

⁵⁸ Randazzo, Chiara, Federico Dell’Acqua, Ethan Mollick, François Candelon, and Karim R. Lakhani, “Cyborgs, Centaurs and Self-Automators,” Harvard Business School Working Paper No. 26-036, 2025.

interaction with generative AI are associated with perceptions of poorer collaborative problem-solving outcomes,⁵⁹ and expert evaluators of AI-supported work penalize outputs less when they are told that workers retained strategic control over decisions about work content and outcomes.⁶⁰

Generative AI may be particularly well-suited to supporting collective efficacy because it can help individuals integrate performance-relevant information from the team level into individual decision making. A recent large-scale field experiment found that “teams paired with AI produced outputs that better integrated diverse perspectives than teams without AI.”⁶¹ Simulations of AI teammates configured with shared goals and social perceptiveness have been shown to improve team-level collective intelligence by 11 to 16 percentage points over human-only baselines.⁶² These empirical findings support theoretical work identifying mechanisms of AI’s role in supporting team performance, including the development of shared mental models (also known as Theory of Mind) and shared goal alignment.^{63,64} Taken together, AI systems that expand user access to shared context (domain expertise, mental models, shared goals, etc.) can enhance collective efficacy.

Indeed, the authors’ earlier work on *vibe teaming*—an experimental approach to human-AI collaboration that positions AI tools to elevate human-to-human interaction—has shown that integrating rich human context into standard proprietary AI interfaces (using transcripts of structured team deliberations designed to elicit diversity and novelty paired with context-calibrated models and iterative drafting techniques) is associated with notable increases in efficiency and quality of outputs.⁶⁵ These techniques helped increase human-to-human interaction time as a portion of total work time while also

⁵⁹ Gaoxia Zhu, Vidya Sudarshan, Jason Fok Kow, and Yew Soon Ong, “Human-Generative AI Collaborative Problem Solving: Who Leads and How Students Perceive the Interactions,” arXiv, May 19, 2024, <https://arxiv.org/abs/2405.13048>.

⁶⁰ Jin Kim et al., “People Reduce Workers’ Compensation for Using Artificial Intelligence (AI),” preprint, January 22, 2025, arXiv:2501.13228, <https://arxiv.org/abs/2501.13228>.

⁶¹ Fabrizio Dell’Acqua et al., “The Cybernetic Teammate: A Field Experiment on Generative AI Reshaping Teamwork and Expertise,” Harvard Business School Strategy Unit Working Paper No. 25-043 (March 28, 2025), SSRN, <https://ssrn.com/abstract=5188231>.

⁶² Samuel Westby and Christoph Riedl, “Collective Intelligence in Human-AI Teams: A Bayesian Theory of Mind Approach,” in Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI-23), 2023, <https://arxiv.org/abs/2208.11660>.

⁶³ Rafael Kaufmann, Pranav Gupta, and Jacob Taylor, “An active inference model of collective intelligence,” *Entropy* 23, no. 7 (2021): 830, <https://doi.org/10.3390/e23070830>.

⁶⁴ Michael S. Harré, Catherine Drysdale, and Jaime Ruiz-Serra, “Theory of Mind Enhances Collective Intelligence,” arXiv preprint arXiv:2411.09168, November 14, 2024, <https://arxiv.org/abs/2411.09168>.

⁶⁵ Jacob Taylor and Kershlin Krishna, “Vibe teaming: How human-human-AI collaboration could disrupt knowledge work for the world’s toughest challenges,” *Global Economy and Development*, Working Paper 193 (Washington, D.C.: Brookings Institution, June 2025).

increasing perceived work enjoyment and perceptions of flow—a psychological (and neurocognitive) state associated with mastery.⁶⁶

This emerging pattern—where increased user-control over context within AI deployment environments supports more productive engagement with AI—parallels research on technology deployments in larger institutional settings like education, where implementation support (dedicated lab coordinators, protected practice time, supervised use) increases engagement and efficacy.⁶⁷ Taken together, existing evidence suggests outcomes of generative AI use—including efficacy and mastery—could be mediated by the level of user control over context within the AI deployment environment.

The limitation of this evidence is that it emerges from experiments conducted within default AI deployment paradigms. The question of whether a fundamentally different deployment environment could support different outcomes for cognitive agency has, until very recently, remained speculative. In early 2026, that question moved from theoretical to practical.

The OpenClaw movement

In January 2026, an open-source project called OpenClaw went viral, surpassing 100,000 GitHub stars within weeks.⁶⁸ OpenClaw is billed as an agent harness for creating LLM-based AI agents: software that enables users to interact with LLMs and direct their activity across any other software applications accessible via an application programming interface (API). But under the principles of context and context engineering reviewed above, OpenClaw and other similar open-source software projects can also be more fundamentally understood as context harnesses: digital infrastructure that places the curation and control of context—domain knowledge, reasoning patterns, expertise, and memory—in users' hands, in human-readable files they own and can share, stored on hardware that they own or control.

Leading AI companies' response to the OpenClaw moment was swift and revealing. In February 2026, OpenAI hired OpenClaw's lead developer to lead next-generation agent development.⁶⁹ Meta acquired Moltbook in March, a parallel project enabling decentralized agent-to-agent interaction, bringing its founders into Meta

⁶⁶ Kotler, Steven, Michael Mannino, Scott Kelso, and Richard Huskey. "First Few Seconds for Flow: A Comprehensive Proposal of the Neurobiology and Neurodynamics of State Onset." *Neuroscience & Biobehavioral Reviews* 143 (2022): 104956. <https://doi.org/10.1016/j.neubiorev.2022.104956>.

⁶⁷ Oreopoulos, Philip, Oliver Keyes-Kryszakowski, and Deepak Agarwal. How In-School Supervised Ed-Tech Support Produces Massive Learning Gains: A Khan Academy Field Experiment in India. NBER Working Paper 34683, January 2026. <https://www.nber.org/papers/w34683>

⁶⁸ Andrea, "OpenClaw Smashes Records: The Viral AI Agent Is Shaking Up GitHub," Hyperright, January 30, 2026, <https://hyperright.com/openclaw-ai-assistant-rebrand/>.

⁶⁹ Gyana Swain, "OpenAI Hires OpenClaw Founder as AI Agent Race Intensifies," InfoWorld, February 16, 2026.

Superintelligence Labs.⁷⁰ Anthropic, building on the success of Claude Code, accelerated its own shift toward distributed agentic products (like Claude Co-Work) before restricting use of third-party integrations like OpenClaw in April.⁷¹ Taken together, these moves signaled a shared understanding across the industry that the competitive frontier extends beyond model capability alone—now largely commoditized across frontier developers—to controlling how human context is captured and utilized.

But OpenClaw also demonstrated something the industry response only underscored: the possibility of a highly accessible alternate deployment architecture in which users can exercise much greater control over their context. Technical expertise has served as a traditional barrier to user control over interactions within consumer and enterprise software applications. Control over the interaction environment was ceded to software developers and firms, not because users did not want it, but because exercising it required developer time, expertise, and resources they did not have.⁷² The underlying software was something the median user could neither build nor modify.⁷³

As OpenClaw has powerfully shown, generative AI potentially disturbs this status quo. Because LLMs can act as a kind of proxy software developer—configurable in natural language, capable of writing and live-editing the code that sits between a user and the digital tools they rely on—the work of shaping the interaction environment is moving within reach of people who have never written code. OpenClaw’s uptake signals a democratization of software development for diverse and non-technical users interested in building local-scale, personalized digital environments for interaction with LLMs. At the time of writing, OpenClaw and equivalent open-source harnesses boast an impressive global and sectoral footprint. OpenClaw has amassed over 346,000 GitHub stars (making it one of the most-starred projects on GitHub), roughly 38 million website visitors, and 3.2 million users.⁷⁴

Beyond its significance as a bottom-up groundswell of user demand for greater control over AI deployment, the OpenClaw moment is significant for the public policy responses it has received. In China, Tencent and Baidu have amplified diffusion by embedding OpenClaw directly into WeChat and Baidu’s flagship search app, exposing agentic workflows to hundreds of millions of existing users and driving early adoption in finance,

⁷⁰ Reuters, “Meta Acquires AI Agent Social Network Moltbook,” March 10, 2026.

⁷¹ Andrea, “Anthropic Shifts API Strategy for OpenClaw, Ending Pro Subscription Access for Third-Party Agents,” Hyperight, April 6, 2026.

⁷² Jean Tirole, “Competition and the Industrial Challenge for the Digital Age,” background paper for the IFS Deaton Review, April 3, 2020, https://www.tse-fr.eu/sites/default/files/TSE/documents/doc/by/tirole/competition_and_the_industrial_challenge_april_3_2020.pdf.

⁷³ Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (New York: PublicAffairs, 2019).

⁷⁴ OpenClaw VPS. “OpenClaw Statistics 2026: Growth, Users, Security, Data.” OpenClaw VPS Blog, April 2, 2026. <https://openclawvps.io/blog/openclaw-statistics>.

internet platforms, and manufacturing and supply-chain firms.⁷⁵ Local governments’ “lobster policies” in Shenzhen and Wuxi signal official support for open, locally deployable agent infrastructure as a way for firms, especially small and medium-sized enterprises, to “own the digital means of production.”⁷⁶

Considered holistically, the OpenClaw movement marks the emergence of a fundamentally different paradigm of personal computing that could, over time—and thanks to AI itself—help center human control over context in their interactions with AI through self-managed software, hardware, and related competencies. If web 2.0 lowered the cost of communication and commerce while keeping platform software inside the platform, generative AI is lowering the cost of building the platform itself—and with it, the long-standing objection that users cannot realistically control the environments in which they collaborate with digital systems. This is at least true for the data, memory, and context layers of generative AI deployment infrastructure, even as separate questions remain about how increasingly powerful foundation models and autonomous agentic systems should be safely evaluated, monitored, and constrained within this infrastructure.⁷⁷

Building as policymaking

How should policymakers respond to these events? OpenClaw’s emergence creates new opportunities to research the relationship between variation in generative AI deployment infrastructure and cognitive agency across individuals, collectives, and societies. But for this paper to simply call for more research on this topic would fail to grasp the significance of the transformation in motion with generative AI generally and OpenClaw more specifically.

A more apt response might be for policymakers to join the OpenClaw movement and start building. A growing field of scholarship recognizes the need to shift policymaking culture and practice from evidence analysis and synthesis toward experimental prototyping,

⁷⁵ Reuters, “Tencent Integrates WeChat with OpenClaw AI Agent amid China Tech Battle,” Reuters, March 22, 2026, updated March 23, 2026, <https://www.reuters.com/technology/tencent-integrates-wechat-with-openclaw-ai-agent-amid-china-tech-battle-2026-03-22/>; CNBC, “China’s Baidu Adds OpenClaw AI into Search App for 700 Million Users ahead of Lunar New Year,” CNBC, February 13, 2026, <https://www.cnbc.com/2026/02/13/baidu-openclaw-ai-search-app-integration-china-lunar-new-year.html>.

⁷⁶ Cang Wei, “Wuxi High-Tech Zone Launches Supportive Policies for OpenClaw Project,” China Daily, March 10, 2026, <https://www.chinadaily.com.cn/a/202603/10/WS69afd0f6a310d6866eb3d021.html>; Tempering these more official signals, Chinese AI policy [analysts](#) and industry-native [critics](#) warn that OpenClaw’s apparent scale may reflect labor-market anxiety, political signaling, and hype-driven, fragile deployments as much as the promise of durable productivity gains or returns to human capital.

⁷⁷ Li, Miles Q., and Benjamin C. M. Fung. “Security Concerns for Large Language Models: A Survey.” Preprint, McGill University, 2025.

rapid iteration, and networked learning.^{78,79} Policymaking that does not participate in this dynamic learning loop fails to address what it is trying to govern—and this mismatch rises sharply for digital technologies that change faster than policy cycles can.^{80,81}

Generative AI is exciting because AI changes how AI policy research can usefully participate in more agile forms of policy deployment.⁸² It enables a form of prototyping as policymaking—standing up an instance of a policy argument as a model of its larger potential form as a means to test it. As Ethan Mollick has observed, generative AI has placed everyone working with these systems in a de facto R&D role: experimenting with prompts, discovering workflows, and developing tacit knowledge about what works.⁸³ A new challenge for policy research is to make this local AI experimentation public and visible—codifying emerging practices, testing them systematically, and feeding results back into a shared evidence base to improve public policymaking and societal outcomes.⁸⁴

In this spirit, this paper aims to codify and elevate common elements emerging from a growing community of users (which includes the authors) who are benefiting from increased control over context in their interactions with generative AI afforded by a new computing paradigm.

Introducing context-maxxing

A growing community of lead users—spanning research studios, software engineering startups, investment teams, and independent analysts—is already converging on recognizable patterns for operationalizing this new computing paradigm for interacting with generative AI. Much of this experimentation is visible in public repositories and

⁷⁸ Doleac, Jennifer, and Anna Harvey, with Jonathon Attridge, Erin Dalton, Dan Kreisman, Weston Merrick, Anthony F. Pipa, Jenni Owen, and Jim Sullivan. “Building State and Local Government Innovation Capacity: Investing in University–Government Innovation Partnerships.” 17 Rooms, Brookings Institution and Rockefeller Foundation, May 16, 2025. <https://www.brookings.edu/articles/building-state-and-local-government-innovation-capacity-investing-in-university-government-innovation-partnerships/>.

⁷⁹ Jennifer Pahlka, *Recoding America: Why Government Is Failing in the Digital Age and How We Can Do Better* (New York: Metropolitan Books, 2023)

⁸⁰ Beth Simone Noveck, *Solving Public Problems: A Practical Guide to Fix Our Government and Change Our World* (New Haven, CT: Yale University Press, 2021).

⁸¹ Mike Bracken, David Eaves, and Michelle Wronski, *The Service Gap: Europe’s International Digital Strategy 2025* (Brussels: The Lisbon Council, 2025), <https://lisboncouncil.net/the-service-gap-europe-international-digital-strategy-2025/>.

⁸² Taylor, Jacob, and Scott E. Page. “AI Is Changing the Physics of Collective Intelligence—How Do We Respond?” Brookings Institution, December 16, 2025. <https://www.brookings.edu/articles/ai-is-changing-the-physics-of-collective-intelligence-how-do-we-respond/>.

⁸³ Ethan Mollick, “Latent Expertise: Everyone Is in R&D,” *One Useful Thing*, June 20, 2024,

⁸⁴ This approach does not replace the need for traditional (top-down) regulatory frameworks, standards, and public investment because without top-down signals of support, bottom-up movements risk being foreclosed by incumbent interests, as Anthropic, Meta, and OpenAI’s rapid responses to OpenClaw have already demonstrated. This paper aims to contribute to conditions where both modes reinforce each other: Bottom-up prototyping generates evidence and demonstrates feasibility, while top-down interventions provide infrastructure, legitimacy, and guardrails that allow experimentation to scale.

practitioner communities; another portion is happening inside startups with strong incentives to keep their methods proprietary. While many users appear to be using OpenClaw to automate narrow or specific tasks, others are beginning to explore ways in which open-source harnesses provide critical infrastructure for holistically managing context across digital interactions, workspaces, and knowledge bases.

This paper refers to this emerging approach as context-maxxing, or using self-managed hardware and software to maximally control user-generated context to support interactions with generative AI.

The “-maxxing” suffix—an internet meme shorthand for obsessive, algorithmically mediated optimization⁸⁵—is used deliberately in the spirit of reclaiming the meme in service of human interests. Tokenmaxxing has emerged as a dominant management practice inside firms deploying generative AI, optimizing for maximum employee token consumption; an internal Meta leaderboard briefly went viral in early 2026 for ranking employees by individual token usage. “Outcomemaxxing” offers a sensible counter, on the logic that firms should optimize for what AI produces rather than how much of it is consumed.⁸⁶

Crucially, neither side asks who controls the context through which those outputs are produced, or what this algorithmic optimization logic does to the cognition of the people subject to it. Context-maxxing enters this debate by centering on whether the human-provided context is being selected and structured in ways that support cognitive agency—user control, efficacy, and mastery—with generative AI.

This paper codifies two preliminary, mutually reinforcing components of context-maxxing:

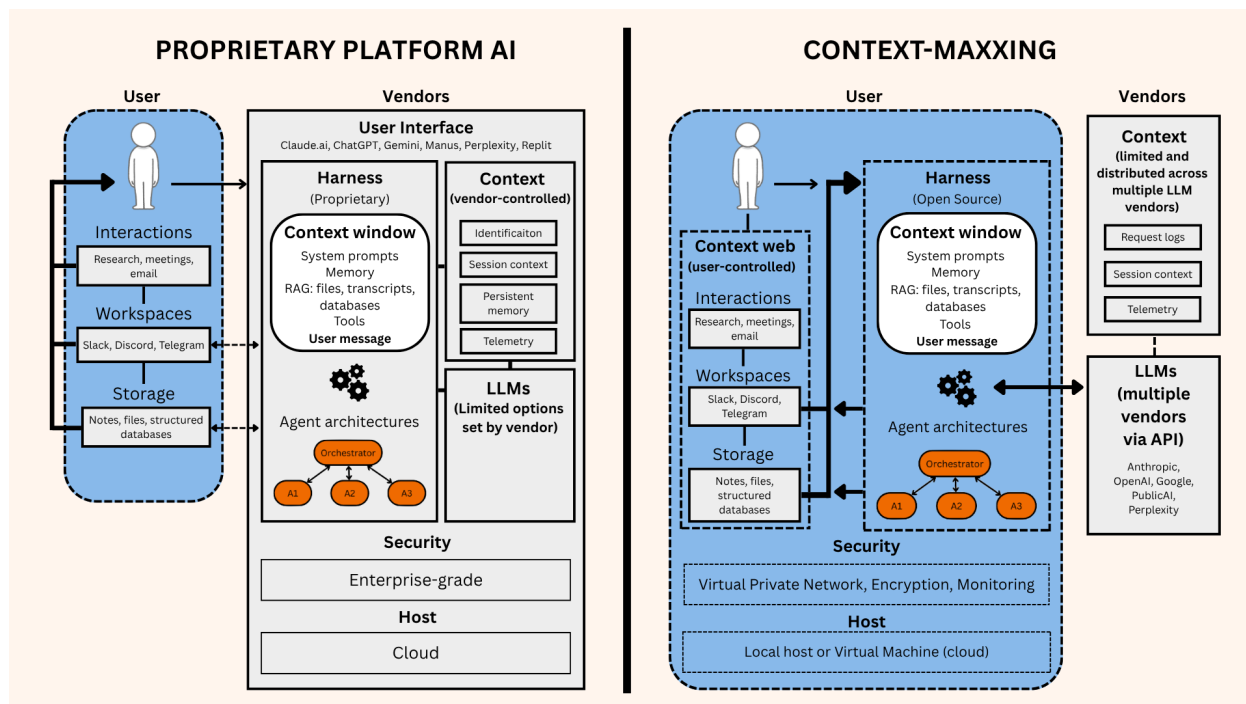
1. **User-controlled digital infrastructure:** Five building blocks sit under user control rather than behind a single vendor wall: (1) an open-source agent harness (e.g., OpenClaw); (2) LLM access via API to multiple providers; (3) a user-assembled context web that holds domain knowledge, reasoning patterns (recurring analytical

⁸⁵ Walker, Aidan. “Clavicular and Contentmaxxing: The Next Step After Groyperfication.” *How to Do Things With Memes*, January 20, 2026. <https://howtodothingswithmemes.substack.com/p/clavicular-and-contentmaxxing>.

⁸⁶ Bousquette, Isabelle. “Why Some Companies Say AI ‘Tokenmaxxing’ Is Key to Survival.” *Wall Street Journal*, April 14, 2026. Tokenmaxxing refers to a dominant management practice emerging inside many firms deploying generative AI, of optimizing for maximum employee token consumption. An internal Meta leaderboard briefly went viral in early 2026 for ranking employees by individual token usage and awarding titles like “Token Legend.” Reporting has documented a broader industry debate in which advocates defend token-maxxing as existential, while critics propose outcome-maxxing as a more disciplined alternative, optimizing for what the AI produces rather than how much of it is consumed. Notably, both sides argue within the same frame. Neither asks who controls the context through which those outputs are produced. Neither asks what leaderboard-driven AI consumption does to the cognition of the person being ranked. Existing debates are internal to the proprietary paradigm they critique.

frameworks and decision rules), and memory in human-readable files; (4) security protocols (least-privilege tool access, credential management, sandboxing, audit logs monitored by agents and humans, and human approval gates); and (5) persistent hosting (local hardware, cloud virtual machines, or hybrid deployments).⁸⁷

Figure 2. Two generative AI deployment infrastructures in 2026

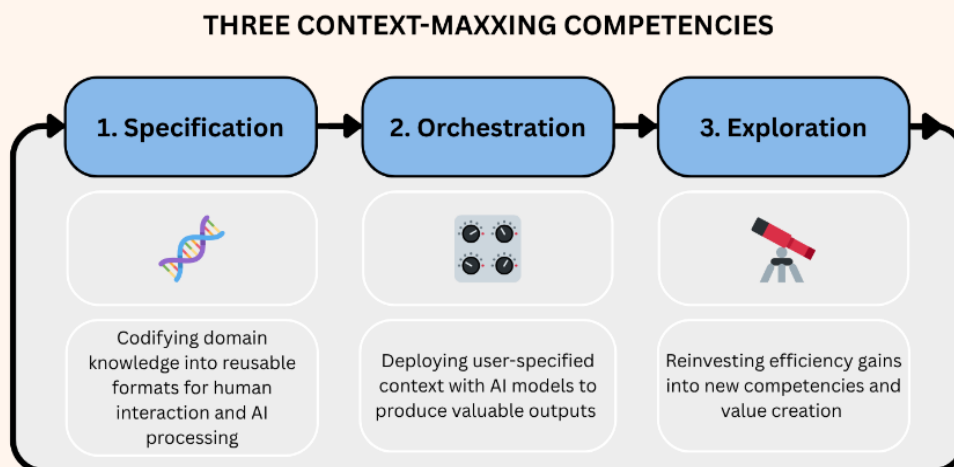


Note: Figure depicts proprietary platform user interfaces (left) versus context-maxxing (right). In context-maxxing, an open-source harness allows users to more meaningfully control and accumulate context within a self-managed computing environment (blue space). Proprietary platform user interfaces (grey space) typically limit user control over context by default or by design. Dotted-line boxes denote features of the AI deployment infrastructure that users can reliably control. Arrows denote flow of information between infrastructure building blocks. Proprietary platform user interfaces increasingly offer direct integration with user-controlled applications (e.g., email, storage, calendar), denoted by the dotted arrow lines.

Source: Visualization by authors.

2. **Emerging competencies that support cognitive efficacy and mastery:** This paper provisionally identifies specification (codifying domain knowledge into reusable context assets), orchestration (operationalizing and optimizing context through human-AI workflows), and exploration (reinvesting efficiency gains into new frontiers of value creation) as necessary preconditions for value creation and shared problem solving with AI.

⁸⁷ These security measures do not eliminate all privacy risks: unless a local language model is used, selected context will still be transmitted to the model provider at inference time. The relevant gain is therefore more intentional user control over what context is sent, to which provider, under what terms, and with what permissions.

Figure 3. Three reinforcing competencies of context-maxxing

Source: Visualization by authors.

A practical playbook

To this end Part 2 of this paper codifies emerging context-maxxing experiments and practices into a playbook for policy and decisionmakers working on or with generative AI, including those interested in supporting user cognitive agency with generative AI as a path to new forms of value creation and shared problem-solving.

The authors include some of their own learnings in this playbook, having implemented and iteratively refined user-controlled context engineering environments at Brookings.⁸⁸ These implementations operationalized the five infrastructure building blocks and three competencies described below across several successive multistakeholder policy convenings, providing a practitioner perspective that informs the playbook alongside observations from a broader community of practice. While not a controlled experiment, these experiences are offered in the spirit of policymaking as prototyping—a grounded account of what context-maxxing looks like in judgment-dependent policy work—and what it demands of practitioners who attempt it.

Directions for future research and public policy implications are discussed in Part 3.

⁸⁸ A forthcoming paper by the authors will describe the Brookings implementation in greater detail and present findings from its application.

Table 1. Dimensions of cognitive agency under default proprietary deployment vs. context-maxxing

| Cognitive agency dimension | Default proprietary AI For example: ChatGPT, Claude.ai, Gemini, Perplexity, Manus, and Replit | Context-maxxing (Open-source harness, LLMs via API, and user-controlled context web, security, and hosting) |
|-----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Control | <ul style="list-style-type: none"> User context is stored and operationalized inside the vendor platform; retention, export, reuse, and any use for product or model improvement are governed by product tier, vendor settings, and terms of service. Users can choose from a selection of LLMs dictated by the vendor. Computational costs are typically abstracted through subscription-based payment. Users can configure their user interface to optimize context (e.g., defining system prompts, creating project folders, and choosing style preferences). Users can integrate select applications (e.g., email, calendar, messenger, and file storage) following vendor-defined protocols | <ul style="list-style-type: none"> User context, model outputs, and agent interactions are held in readable and editable files on user-managed hardware. Users can choose between multiple models from different vendors; cost, latency, and privacy tradeoffs are explicit thanks to API model use, driving accountability and responsibility at the task level. User-selected context is sent to the model provider through API calls; vendor retention and reuse are governed by API terms, enterprise controls, and any available zero-retention settings. Users can integrate any application accessible via an API. |
| Efficacy | <ul style="list-style-type: none"> Efficacy depends on a user’s ability to assemble relevant context and operationalize workflows within the specific proprietary interface. Ability to share and accumulate context between users (e.g., between teammates) is defined by platform features. | <ul style="list-style-type: none"> Efficacy depends on the quality of user-defined (and reusable) context and structured workflows. Ability to share and accumulate context between users is defined by users and depends on permissions and access to shared workspaces or common file storage. |
| Mastery | <ul style="list-style-type: none"> Context and capabilities remain partly platform-entangled: artifacts may be exportable, but the accumulated working context, memory, and workflow patterns are difficult to reuse elsewhere. | <ul style="list-style-type: none"> Capability developed through context-maxxing is retained within a user-controlled computing environment. Specification assets and context webs are portable by construction, accumulate and potentially appreciate through use; they travel with the practitioner. |

Note: This table is a non-exhaustive comparison focused on cognitive agency rather than a comprehensive ranking of deployment models. Proprietary AI products may offer advantages in reliability, agentic capability, cost predictability, ease of setup, and safety or incident-response infrastructure. Context-maxxing, by contrast, can involve higher API costs, greater technical setup, ongoing operational burdens, and additional user responsibility for security and oversight. The point of comparison is therefore not that context-maxxing is strictly dominant, but that it shifts more control over context, workflow design, portability, and verification infrastructure toward the user or team.

Part 2. A playbook for context-maxxing

Context-maxxing can be understood as the practical discipline of maximizing user control over human-generated context to improve the quality of human–AI collaboration. The claim is that context-maxxing could—relative to proprietary platform AI deployments—help preserve and expand human cognitive agency with generative AI.

Across an emerging community of frontier users building and deploying personalized context engineering systems—spanning research studios,⁸⁹ software engineering startups,⁹⁰ and investment teams⁹¹—common context-maxxing infrastructures and competencies are beginning to emerge. This playbook describes these elements with a view to understanding how they might support cognitive agency with generative AI.

The playbook likely only captures one part of a broader wave of experimentation now forming around open-source agent harnesses. Some users are applying these tools narrowly to automate discrete tasks, such as monitoring calendars, scanning emails, or triggering scheduled actions. Others are building directly on the harness software itself, publishing new modules, integrations, and workflows through open-source repositories.⁹² Indeed, several early users are also reporting skepticism about the overall value of OpenClaw-style deployments given their significant initial startup costs relative to demonstrable immediate impact on user productivity or creativity.⁹³

A further set of advanced early users appears to be pushing toward a more holistic approach that treats harnesses as infrastructure for managing context across interactions, knowledge bases, and integrations with external applications (or tools) as a basis for optimizing business processes and developing outputs. Much of this experimentation is visible in public repositories and communications among practitioner communities (including GitHub, Reddit, and newsletters), while another portion is likely happening quietly inside startups, software teams, and firms that have strong incentives to keep their methods proprietary. This paper distills some recurring elements from this activity into a practical framework for understanding how user-controlled deployment

⁸⁹ Azhar, Azeem. "Behind the Scenes of My AI Agent." Exponential View (podcast), February 27, 2026. <https://www.exponentialview.co/p/behind-the-scenes-of-my-ai-agent>.

⁹⁰ HumanLayer, "HumanLayer - Close Your Editor Forever," accessed April 15, 2026, <https://www.humanlayer.dev/>

⁹¹ J.P. Morgan, "Trading Insights: AI in the Macro Process, with Balyasny's Head of Macro Research & Chief Economist," Market Matters podcast, December 13, 2024, <https://www.jpmorgan.com/insights/podcast-hub/market-matters/ai-in-the-macro-process-balyasny>

⁹² OpenClaw. "openclaw/openclaw." GitHub repository, created January 29, 2026. <https://github.com/openclaw/openclaw>.

⁹³ Herrman, John. "My Adventures With 'The AI That Actually Does Things'." New York Magazine, *Intelligencer*, April 28, 2026. <https://nymag.com/intelligencer/article/my-adventures-setting-up-openclaw-agent.html>.

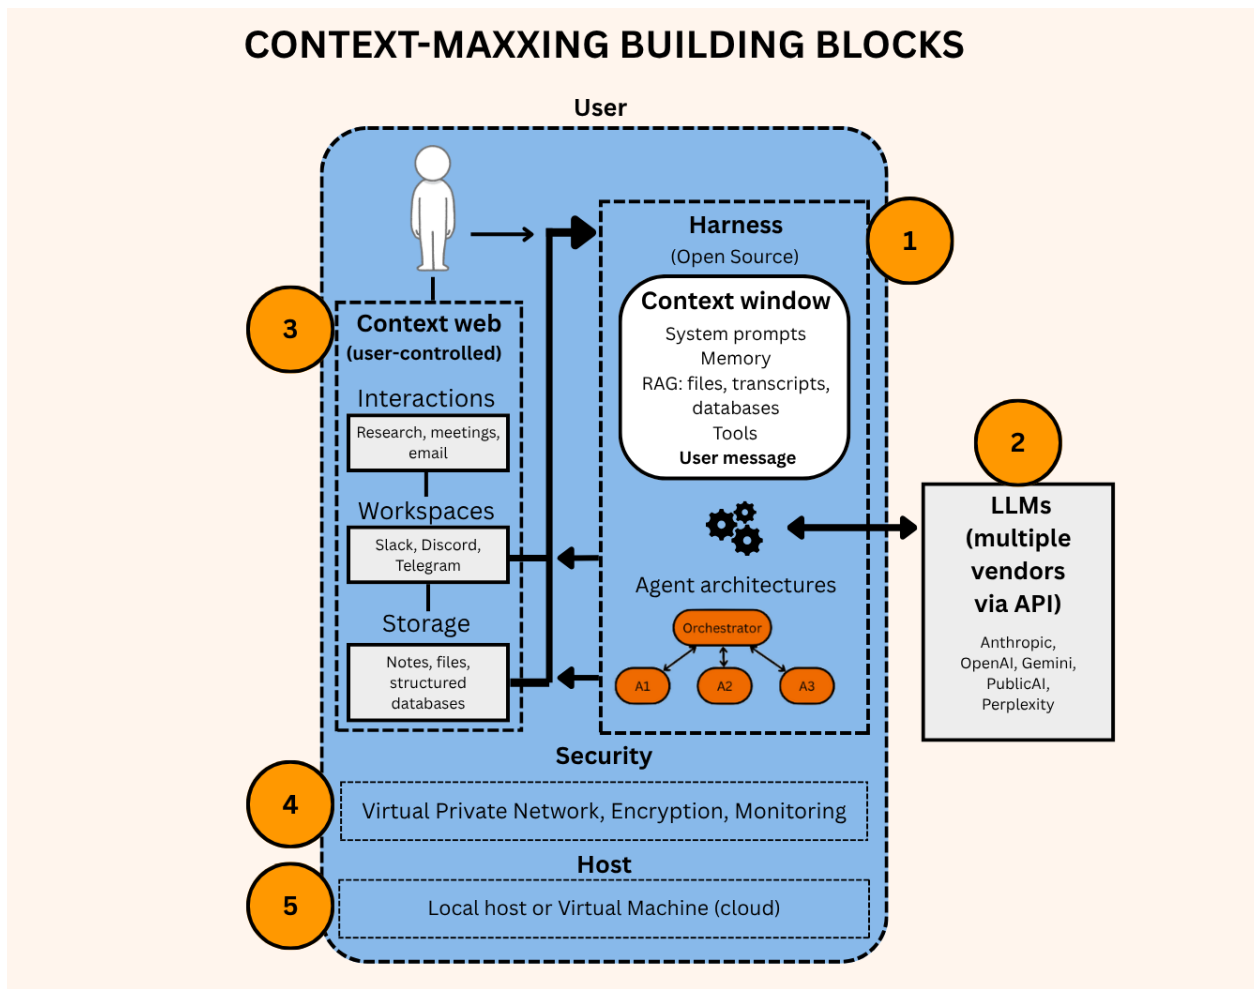
infrastructure and emerging context-maxxing competencies can support cognitive agency with generative AI.

Specifically, the sections below present a snapshot of an evolving set of (i) five common digital infrastructure building blocks (open-source agent harness, LLM access, context web, security hardening, and persistent hosting) and (ii) three reinforcing competencies (specification, orchestration, and exploration).

2.1 Five infrastructure building blocks

A full generic implementation of common digital infrastructure building blocks of context-maxxing is represented in Figure 4.

Figure 4. Five emerging infrastructure building blocks for context-maxxing



Source: Visualization by authors.

1. Harness

An agent harness (or context harness, as reframed above) is software that lets users connect AI models to selected files, instructions, memory, and tools. Instead of accessing LLMs through a proprietary user interface platform like Claude.ai, ChatGPT, or Perplexity, the harness provides an environment users can configure themselves that mediates interaction between LLMs and a defined operating context. When the harness is open-source and self-hosted—on a personal machine or managed server—practitioners gain greater control over how the AI system is used because the operating context lives in plain-text files they can read and edit rather than being stored inside an opaque proprietary platform.⁹⁴

Once running, the harness can be connected to any chat interface that allows software connections through an API, such as Discord, Slack, or Telegram, enabling user interaction with the same underlying harness across multiple channels. Tool and data-source integrations increasingly default to the Model Context Protocol (MCP), an open standard for connecting AI applications to external systems such as local files, databases, tools, and workflows. This makes the harness an integration layer as well as an interface layer: It governs which parts of the user’s context web are exposed to the model, under what permissions, and in what form.

The harness loads plain-text configuration files that define the user’s agent’s identity, working memory, operating mode, task-handling conventions, curated domain knowledge, and workflow integrations (these files are typically stored in markdown or JSON format and have names like “SOUL.md”, “INSTRUCTIONS.md”, or “MEMORY.md”). Because these context files are human-readable documents under user control, they are fully inspectable, auditable, reusable, and portable (interoperable with other interactions with LLMs beyond a specific implementation). Users can manage these files through version control, track changes over time, clone them to create new harnesses, and use them across software systems. Proprietary AI user interfaces, by contrast, typically allow users some ability to inspect and configure these files (or their equivalents), either by a user directly querying this information through a chat interface, or through settings controlled by the provider.

While OpenClaw is the leading example of a model harness, a growing ecosystem offers alternatives focused on security, efficiency, performance, and research (Appendix A

⁹⁴ Users can install an agent harness in the terminal of their personal computer using a small number of commands that automate dependency management, configuration, and service setup. The [OpenClaw installation](#) is indicative of a new generation of low-code software installers, which increasingly expose natural-language prompts and AI-guided installation wizards, allowing users to complete setup by answering straightforward questions instead of writing configuration code. One-click set ups for OpenClaw is also increasingly available through [cloud hosting providers](#).

Table A1). For policymakers, the crucial takeaway is that whoever controls the harness controls the context.

By placing configuration in editable plain-text files under user ownership, the harness directly supports the control dimension of cognitive agency. Users can examine, modify, and audit the informational environment that shapes their AI interactions—capabilities that are structurally limited in proprietary interfaces where equivalent configuration is stored inside the platform. Because context files are human-readable and stored locally, they can travel with the user rather than remaining entangled with a specific vendor.

2. LLM access via API

LLMs are integrated into a harness through an API key, which allows the harness to send selected inputs to the model provider (or a locally hosted LLM) and receive outputs. As examined in Part 1, model performance—in both proprietary interfaces and open-source agent systems—depends on what context is made available to the model at inference time. In proprietary systems, users exert some control over how their context is used and stored (e.g., through what they choose to include in the context window of each query, the creation of “project folders,” and integrations with external applications), with parameters ultimately dictated by the vendor. In an open-source architecture, by contrast, the user actively determines exactly what retrieved memory, files, tool calls, and prior exchanges to include in each model call. This can give users greater control over context assembly, make that process more transparent and customizable, and reduce reliance on proprietary platforms to store and manage the contextual assets that shape model performance.

LLM interactions through an API are billed per token, making design choices (history, tool description verbosity, summarization frequency) critical for cost, latency, and performance. Compared to traditional flat-rate subscription services of proprietary vendor platforms, the close coupling between context specification and orchestration and AI utilization costs establishes a real-time feedback loop where (in)efficiencies of context engineering are immediately reflected in expenditures. This compels users to more carefully consider the selection, structure, and computational demands of the input context.⁹⁵

Leading users utilize a tiered model approach: frontier models for high-stakes reasoning and complex tool orchestration, mid-tier models for most workflows, and cost-optimized models for bulk tasks. Competitive open-weight models accessed via APIs have narrowed the performance gap substantially, even as the gap fluctuates when new

⁹⁵ There are already signs that the tokenmaxxing’s “More! More! More!” ethos—reflecting an end-2025 era context of cheap and abundant compute—is [shifting](#) to a more prudent stance in 2026 as token budgets increase as a total share of firm operating budgets, compute shortages threaten supply, and as growing decisionmaker demand for evidence of AI’s impact on productivity.

closed-weight frontier systems are released.⁹⁶ The harness's routing logic makes performance-cost-privacy tradeoffs explicit and adjustable (it is easy to dictate model preferences and selection and track token usage per task). This gives users control and choice in a fast-moving and competitive environment of rapid model capability advancements and pricing and privacy shifts, practitioners can rebalance without needing to reorganize their entire digital environment for a new vendor system.⁹⁷

Considered from the perspective of cognitive agency, per-token billing and multi-provider routing support control by making cost, latency, and privacy tradeoffs explicit and adjustable at the task level—a form of transparency and responsibility absent from subscription-based platforms where model selection is limited and computational costs are abstracted. This transparency may over time support cognitive efficacy and performance: The close coupling between context engineering choices and expenditures creates a real-time feedback loop that compels practitioners to optimize the selection, structure, and scope of input context, potentially improving output quality as a byproduct of cost discipline.

3. Context web

A harness's value is defined by the digital systems that LLMs connected to it can navigate to collect and assemble context. This paper understands this distributed space as a user's context web—or the universe of digital information available across a user's digital surfaces. This term emphasizes that valuable context exists across multiple digital surfaces and users gain value from integrating these surfaces to identify and develop relationships among information stored across them.

In practice, a context web spans three layers (see Figure 4 above and Table 2 below): the interaction layer (how new information enters a user or team's context environment—web browsers, email, external and internal meetings), a user's primary workspace and the communication layer (messaging apps, calendars, CRM, project trackers); and a knowledge layer (notes, documents, personal knowledge management tools, file storage, and structured databases). Maximizing available user context means maximizing integration of the right context across user applications across all three layers. For this reason, context-maxxing privileges applications and tools that are open (have an open API) and interoperable (codify information in standard formats like markdown).

By integrating context across interaction, workspace, and knowledge layers, the context web can support cognitive efficacy through richer, more situated information for

⁹⁶ Stanford Institute for Human-Centered Artificial Intelligence, AI Index Report 2026 (Stanford, CA: Stanford University, 2026), 76.

⁹⁷ Thomas Claburn, "Anthropic Tweaks Timed Usage Limits to Discourage Claude Demand during Peak Hours," *The Register*, March 26, 2026, https://www.theregister.com/2026/03/26/anthropic_tweaks_usage_limits/; Raffaele F. Ciriello and Kathryn Backholer, "OpenAI Will Put Ads in ChatGPT. This Opens a New Door for Dangerous Influence," *The Conversation*, January 23, 2026, <https://doi.org/10.64628/AA.e3pvkfugk>.

interaction with LLMs. Over successive applications, the context web also supports cognitive mastery: Integration of context assets across the three layers—meeting transcripts, templates, structured notes, and documents—reduces the marginal cost of future knowledge production.

Table 2. Context web layers

| Layer | Description | Applications and tools | Relevant context |
|------------------------------------------|-----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------|
| Interaction layer | Where new information and signals enter a user's context. | Web browsers (Google Chrome, Firefox), email (Gmail, Outlook), RSS, social feeds (Twitter, Bluesky), in-person and virtual meetings (captured digitally with transcription tools like Otter.ai). | Capture, transcribe, enrich, and route external inputs into workflows and knowledge layers. |
| Workspace and communication layer | Where work is coordinated and communicated . | Internal messaging and collaboration apps (Slack, Discord, WhatsApp), calendars, CRM and project management tools (e.g., Airtable). | Develop ideas, scope tasks, track commitments, timelines, and collaborators; connect rich context to specific actions and outputs. |
| Knowledge layer | Stored knowledge, artifacts, and preferences. | Personal Knowledge Management (PKM) or notetaking apps (Obsidian, Roam Research, Notion), documents (Google Docs), structured databases (Mother Duck, Neo4j). | Store, link, and retrieve durable knowledge that grounds the agent's reasoning. |

4. Security

Connecting an agent to personal and professional systems creates concrete exposure risks. When a harness can read files, send emails, execute code, and access calendars or CRMs, at least three documented classes of risk are emerging: direct compromise of the harness (remote code execution vulnerabilities), theft or misuse of credentials and API keys, and indirect attacks like prompt injection or memory poisoning that cause the agent to perform harmful actions with otherwise legitimate access.⁹⁸ OpenClaw vulnerabilities have enabled WebSocket hijacking and token theft from a browser session, along with command injection, server-side request forgery, sandbox boundary bypasses, and path

⁹⁸ Jeffrey Brainard, "AI Algorithms Can Become 'Agents of Chaos,'" *Science*, March 23, 2026, <https://www.science.org/content/article/ai-algorithms-can-become-agents-chaos>.

traversal leading to local file exposure.⁹⁹ Analyses of Internet-exposed instances have found large numbers of poorly secured OpenClaw deployments with default settings and no network isolation, effectively turning personal agent setups into remotely accessible admin shells.¹⁰⁰

To guard against these security threats, current guidance boils down to treating harnesses as high-value, networked applications.¹⁰¹ Specific measures include isolating harnesses via containers or virtual machines, using a VPN-protected network, enforcing strong authentication protocols, applying least-privilege permissions at the tool and action level, establishing regular security monitoring procedures, and creating kill switches (see Appendix Table B3 for more details).

Security hardening supports cognitive agency primarily through the control dimension. Practitioners who manage their own exposure risks develop a more accountable relationship with the infrastructure than users of proprietary systems where security decisions are made on their behalf. While accountability imposes real operational burdens, it also means users are compelled to understand, modify, and monitor the security of their system.

5. Hosting

Most fundamentally, all of this software infrastructure needs hardware on which it runs. Options include a user's primary workstation (quick to start but couples the agent to everyday computing and frequently a user's sensitive data), a dedicated local machine (physically controlled, always-on, close proximity to local files), or cloud-hosted virtual machines (continuous uptime, easier remote administration, greater compute elasticity, but less control over data residency). Many practitioners are converging on a hybrid model: dedicated local hardware for personal or sensitive workflows, cloud-based

⁹⁹ National Vulnerability Database, "CVE-2026-25253 Detail," National Institute of Standards and Technology, published February 1, 2026, <https://nvd.nist.gov/vuln/detail/CVE-2026-25253>; GitHub Advisory Database, "OpenClaw Affected by SSRF via Attachment/Media URL Hydration," GitHub, published February 17, 2026, <https://github.com/advisories/GHSA-wfp2-v9c7-fh79>; GitHub Advisory Database, "OpenClaw Is Vulnerable to Path Traversal through Path Validation Bypass," GitHub, published March 26, 2026, <https://github.com/advisories/GHSA-hggm-x7r9-mm7v>.

¹⁰⁰ Silas Cutler, "OpenClaw in the Wild: Mapping the Public Exposure of a Viral AI Assistant," Censys, January 31, 2026, <https://censys.com/blog/openclaw-in-the-wild-mapping-the-public-exposure-of-a-viral-ai-assistant/>; Tom Fosters, "Don't Get Pinched: the OpenClaw Vulnerabilities," Kaspersky Official Blog, February 10, 2026, <https://www.kaspersky.com/blog/openclaw-vulnerabilities-exposed/55263>.

¹⁰¹ FounderCoHo, "OpenClaw: The Malware You Installed on Purpose — A Security Playbook: A Layer-by-Layer Risk Analysis of the Fastest-Growing Open-Source AI Agent in History --- For Users, Builders, and Maintainers," FounderCoHo (Substack), March 14, 2026, <https://foundercoho.substack.com/p/openclaw-the-malware-you-installed>.

hosting for distributed teams.^{102,103,104} Dedicated hardware (Mac mini/Framework Desktop) costs roughly \$600 to \$2,000 (as of April 2026); a cloud VM capable of running a harness continuously costs \$10 to \$40 per month.¹⁰⁵

Hardware ownership or management supports control by keeping data residency decisions in user hands. Over time, hosting choices also support mastery: A dedicated computing environment that accumulates configuration, memory, and context assets across applications functions as a persistent infrastructure for capability development—unlike proprietary interfaces where accumulated working context is difficult to export or reuse.

Implementation at Brookings

As knowledge workers at Brookings, the authors implemented context-maxxing for policy research. The authors each stood up a separate instance of the five building blocks described above (see Appendix Figure A1). Both used OpenClaw harness connected via Slack and Telegram; multi-provider API access with tiered model routing (Anthropic, Gemini, and Apertus accessed via PublicAI); a context web spanning Roam Research (personal knowledge management layer), Slack (workspace and communication layer), Airtable (CRM), and Google Drive (file storage), with meeting transcripts captured via Otter.ai (interaction layer); security hardening including VPN (Tailscale), firewall, least-privilege tool access, and managed API keys; and hosting on dedicated locally-managed or cloud-based hardware (Mac mini and Hetzner).

Each implementation took approximately 20-25 hours of practitioner time to stand up, with roughly \$600 in token costs during the setup period. These implementations have been operationalized across successive policy projects with notable control, efficiency, and mastery gains. The competencies described below draw on this experience alongside the broader community of practice, with their implementation to be described in greater detail in a forthcoming paper.

2.2 Three reinforcing competencies

Prominent lead users of the OpenClaw-style approaches are demonstrating the need for a set of deliberate individual and team practices to operationalize the infrastructure: How

¹⁰² Azeem Azhar, "Why I Changed My Mind about Apple," Exponential View, March 19, 2026,

<https://www.exponentialview.co/p/why-i-changed-my-mind-about-apple>.

¹⁰³ Google Cloud, "Edge hybrid pattern," Cloud Architecture Center, last modified January 23, 2025,

<https://docs.cloud.google.com/architecture/hybrid-multicloud-patterns-and-practices/edge-hybrid-pattern>.

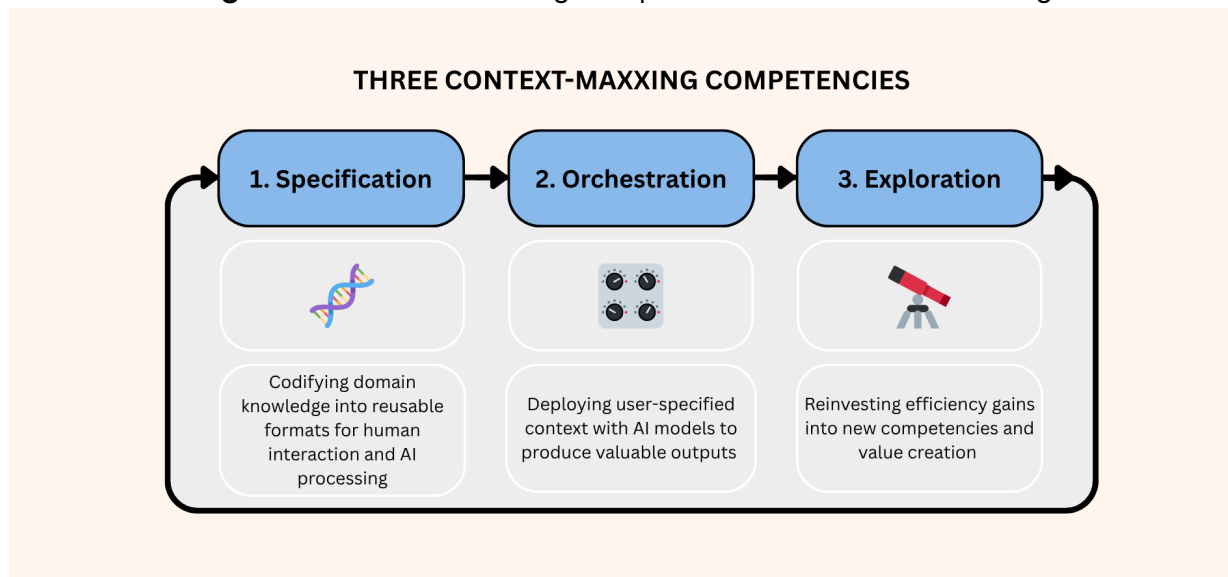
¹⁰⁴ Amazon Web Services, "Edge AI and global inference distribution," AWS Prescriptive Guidance, accessed April 15, 2026, <https://docs.aws.amazon.com/prescriptive-guidance/latest/agent-ai-serverless/edge-ai.html>.

¹⁰⁵ GetDeploying, "VPS Price Comparison," updated March 26, 2026,

<https://getdeploying.com/reference/compute-prices>; Apple, "Buy Mac mini," Apple Store, accessed April 15, 2026, <https://www.apple.com/shop/buy-mac/mac-mini>.

domain expertise is codified into context, how context is routed through human-AI workflows, and how the efficiency gained from AI-enabled work is reinvested into new frontiers of value creation. The paper describes these as three reinforcing competencies: specification, orchestration, and exploration.

Figure 3. Three reinforcing competencies of context-maxxing



Source: Visualization by authors.

Specification: codifying domain knowledge through standardized templates

Specification is the competency of building the informational environment that makes AI useful. It typically involves articulating and codifying formal and tacit knowledge—the professional judgment and domain expertise that typically resides within or between practitioners’ heads—and organizing it in formats and processes optimized for the cognitive constraints and comparative advantages of humans and computers.

This is a non-trivial design challenge. What Richard Sutton terms the bitter lesson suggests that more open-ended computational process automation techniques can often outperform those that are hand-engineered with domain knowledge. An operative question for specification approaches is therefore what to codify and at what level of detail.¹⁰⁶ Frontier users are discovering which slices of their tacit expertise genuinely raise model performance, and which do not. To this end, emerging best practice seeks to balance capturing enough structure to meaningfully guide AI behavior without overwhelming either an LLM’s context window or the cognitive constraints of human reviewers.

Templates

¹⁰⁶ Ethan Mollick, “The Bitter Lesson versus The Garbage Can,” One Useful Thing, July 28, 2025.

One response to this challenge, emerging primarily in software development, is developing standardized templates—or reference architectures—for knowledge transfer between humans and LLMs.¹⁰⁷ Templates are lightweight, immutable, modular structures in which domain knowledge, quality standards, and conventions that shape expert judgment can be structured and re-used. For example:

- Software engineers typically specify architecture conventions, testing standards, and code-review criteria.¹⁰⁸
- A plumbing or HVAC contractor might specify client-property configurations, service-history structures, and the symptom-to-diagnosis-to-repair relationships that make each job a training example for the next.¹⁰⁹
- Lawyers use of the IRAC structure in legal reasoning (Issue, Rule, Application, Conclusion) functions as a standardized template for case work.¹¹⁰
- Policymakers often evaluate the strengths and weaknesses of proposed policymaking innovations through a set of recurring diagnostic questions. In the authors' own implementation, two lightweight templates structure context for human-AI workflows: a "Policy Innovation Template" and a "Policy Critique Template"; see Appendix B).

Templates can help structure knowledge and guide attention to the key information required to research, develop, or produce domain-specific outputs (research memos, feature fixes, cases, code, etc.).¹¹¹

Templates are helpful for context engineering because they are structured enough to shape working memory and attention but flexible enough to apply across varied inputs and outputs, and sized to respect the bandwidth of human minds and context budgets of LLMs. Answers to prompts or questions in a template can be expressed concisely in 50 or 500 words for human readability or deployment to a context window; while also serving as the framework for more in-depth research and development of answers to those same questions. These templates structure accumulation of knowledge around a specific constellation of context that can be used and re-used in interactions with LLMs to develop outputs. The authors have begun referring to the template-developed context

¹⁰⁷ Choro Ulan Uulu et al., "How to Build AI Agents by Augmenting LLMs with Codified Human Expert Domain Knowledge? A Software Engineering Framework," arXiv preprint, January 21, 2026, <https://arxiv.org/abs/2601.15153>.

¹⁰⁸ Matt Tanner, "Enterprise Software Architecture Patterns: The Complete Guide," vFunction, June 18, 2025, <https://vfunction.com/blog/enterprise-software-architecture-patterns/>.

¹⁰⁹ ServiceTitan, "HVAC Diagnostic Chart: What You Need to Know," March 10, 2025, <https://www.servicetitan.com/blog/hvac-diagnostic-chart>.

¹¹⁰ Otto Stockmeyer, "It's All About IRAC," Cooley Law School, accessed April 30, 2026, <https://cooley.edu/blog/its-all-about-irac>.

¹¹¹ Allison Whitten, "New Tool Helps AI and Humans Learn To Code Better," Stanford Institute for Human-Centered Artificial Intelligence, May 1, 2023, <https://hai.stanford.edu/news/new-tool-helps-ai-and-humans-learn-code-better?sf180948227=1>.; The underlying intuition has a parallel in recent work at Stanford CASBS on informational grammars for problem-solving—compact, reusable structures that specify the moves and relationships a reasoning process should be able to make across a class of problems.

assets as “master keys” because of their generalizability and re-use potential for different subject matter and form factors relevant to the team’s policymaking workflows.

Ontologies

As specification matures, practitioners develop ontologies: structured representations of entities, relationships, and temporal dynamics specific to business processes. By making relational knowledge explicit and shared—rather than locked in individual practitioners’ heads—they create a common reference architecture that team members can consult, extend, and correct.

Whereas traditional methods of ontology development in software engineering and machine learning often involve painstaking manual codification and programming—typically for enterprise-scale deployments, generative AI is making ontologies more easily implementable across users’ own digital surfaces, including “Personal Knowledge Management” (or PKM) software like Obsidian or Roam Research (which provide a basic ontology structure of entities, relationships, and time). Lead users of open harnesses are sharing more sophisticated examples using object-oriented approaches and dedicated database software.¹¹²

When considering cognitive agency, specification is the competency most directly associated with cognitive mastery: It requires practitioners to externalize and codify tacit knowledge—a process that deepens domain understanding while creating durable, portable assets. These assets also support efficacy by making domain expertise machine-legible in formats optimized for both human review and LLM context windows, reducing the marginal cost of knowledge transfer to AI systems across successive applications.

Orchestration: operationalizing context through human-AI collaboration workflows

Orchestration describes the emerging competency of how user-specified context is deployed with AI models to produce valuable output or work. The critical design principle is building for context integrity and verifiability—ensuring specification assets are surfaced appropriately, redundant context is excluded, and AI errors are caught before they propagate. Orchestration also means optimizing use of the context window to support model performance and managing performance as a function of token budgets (or compute costs).

¹¹² HEmile, Obsidian Neo4j Graph View, GitHub repository, accessed April 24, 2026, <https://github.com/HEmile/obsidian-neo4j-graph-view>.

Context window optimization

Effective agent orchestration of human-AI workflows requires close management of context-window utilization and token cost. Some practitioners advise keeping active context in a context window (the information provided in a single interaction with an LLM) well below the full limit—often around 40–60% of usable capacity—based on research showing that model performance weakens as attention is diluted across information.^{113,114} In practice, these constraints reward careful model selection, tighter prompt design, and more selective summarization (also known as compaction) and inclusion of information, rather than verbosity.¹¹⁵

Agent architectures

Research on emerging agent architectures indicates that the value of multi-agent coordination is contingent on task structure, not purely on the number of agents.¹¹⁶ This gives rise to a typology of agent architecture design choices that depend on the type of knowledge work and tasks within knowledge work (see Table 3). Parallelizable tasks—analyzing a policy problem from multiple stakeholder perspectives, conducting parallel literature reviews—benefit dramatically from coordination among multiple agents. However, tasks like coding that require one coherent line of thought or sequential reasoning, tend to degrade with multiple sub-agents because coordination overhead costs exceed gains associated with diversity.¹¹⁷ Evidence suggests the one promising way to maintain coherent user context and reduce propagation errors in multi-agent workflows (when an unchecked error that surfaces early in an LLM’s workflow is amplified throughout an entire workflow because it serves as an input to subsequent procedures) is through centralized orchestration: a hub-and-spoke setup where one coordinating (or parent) agent works directly with the user to establish context, assign and delegate work to sub-agents, checks outputs, and synthesize the results for the user.¹¹⁸

¹¹³ Kelly Hong, Anton Troynikov, and Jeff Huber, “Context Rot: How Increasing Input Tokens Impacts LLM Performance,” technical report, Chroma, July 2025, <https://www.trychroma.com/research/context-rot>.

¹¹⁴ Mosh Levy, Alon Jacoby, and Yoav Goldberg, “Same Task, More Tokens: The Impact of Input Length on the Reasoning Performance of Large Language Models,” in Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024), 2024, <https://arxiv.org/abs/2402.14848>.

¹¹⁵ Dex Horthy, “Advanced Context Engineering for Coding Agents,” HumanLayer Blog, August 29, 2025, accessed April 19, 2026, <https://www.humanlayer.dev/blog/advanced-context-engineering>

¹¹⁶ Yubin Kim et al., “Towards a Science of Scaling Agent Systems,” arXiv preprint, December 8, 2025, <https://arxiv.org/abs/2512.08296>.

¹¹⁷ Krishnan, Rohit. “Why Coase Needs Hayek: Sometimes Smart Planners Lose to Simple Markets.” Strange Loop Canon, May 1, 2026. <https://www.strangeloopcanon.com/p/why-smart-planners-lose-to-simple>

¹¹⁸ Ibid. Research on agent architectures shows centralized multi-agent coordination improved performance by 80.9% over single-agent baselines on parallelizable tasks. Sequential reasoning tasks tended to degrade with multiple agents because coordination overhead exceeded distribution gains.

The authors configured structurally similar multi-agent architectures in their implementations at Brookings. A single orchestrator agent routes tasks to dedicated sub-agent teams for research and development (using the Policy Innovation Template) and critique (using the Policy Critique Template), before routing refined outputs to execution sub-agents for specific form factors (meeting briefs, convening tools, public-facing policy memos, or working papers). This centralized architecture aligns with emerging evidence favoring hub-and-spoke coordination for judgment-dependent tasks.

Table 3. Common agent architectures

| Agent structure | |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>Single agent: One agent handles the full task in a unified context stream. This arrangement avoids coordination overhead and may be most suitable for tasks that require holistic or sequential reasoning without decomposition. But this arrangement cannot exploit parallel decomposition when the task is splittable.</p> | <pre> graph TD User((User)) --> A1((A1)) A1 --> Output[Output] </pre> |
| <p>Independent multi-agent system: Multiple agents work separately and combine outputs at the end. Independence maximizes parallelization with minimal coordination, but is especially vulnerable to unchecked error propagation because shared context is not maintained after initial user specification.</p> | <pre> graph TD User((User)) --> A1((A1)) User --> A2((A2)) User --> A3((A3)) A1 --> Aggregation[Aggregation] A2 --> Aggregation A3 --> Aggregation </pre> |
| <p>Decentralized multi-agent system: Agents coordinate through peer-to-peer communication. Suitable for exploration-heavy tasks such as dynamic web navigation, but requires sequential communication to ensure user preferences are maintained.</p> | <pre> graph TD User((User)) --> A1((A1)) User --> A2((A2)) User --> A3((A3)) A1 <--> A2 A2 <--> A3 A1 <--> A3 A1 --> Aggregation[Aggregation] A2 --> Aggregation A3 --> Aggregation </pre> |
| <p>Centralized multi-agent system: A central orchestrator delegates and synthesizes sub-agent work. Research suggests this arrangement is suitable for parallelizable reasoning and error containment.</p> | <pre> graph TD User((User)) --> Orchestrator((Orchestrator)) Orchestrator --> A1((A1)) Orchestrator --> A2((A2)) Orchestrator --> A3((A3)) </pre> |
| <p>Hybrid centralized multi-agent system: An orchestrator plus peer communication combines hierarchy with lateral exchange. Requires higher coordination overheads between agents, which can reduce efficiency relative to simpler designs.</p> | <pre> graph TD User((User)) --> Orchestrator((Orchestrator)) Orchestrator --> A1((A1)) Orchestrator --> A2((A2)) Orchestrator --> A3((A3)) A1 <--> A2 A2 <--> A3 A1 <--> A3 </pre> |

Source: Derived from [Kim et al., 2025](#).

To this end, several lead users appear to be opting for a “Chief of Staff” (or what the authors term a “Chief Context Officer”) that delegates specialized work to sub-agents, reviews intermediate outputs, and reintegrates them into formats suitable for human review.¹¹⁹ This architecture can incorporate critique sub-agents that challenge outputs and flag errors, creating a Socratic dynamic that strengthens final products.¹²⁰ In the authors’ own experience, a Chief of Staff-style orchestrator with global and evolving context for user needs and priorities helps provide a user-friendly interface to more elaborate multi-agent workflows underneath.

Process decomposition

Orchestration also involves process decomposition—structuring complex work for AI engagement into sub-processes. Complex knowledge work typically involves distinct cognitive phases: gathering material, diagnosing problems, designing approaches, executing, and reviewing. When these are presented as a single undifferentiated prompt, the model must manage multiple modes simultaneously within one context window, potentially degrading performance on all. Decomposition keeps individual context windows focused, creates natural checkpoints for human review, and produces intermediate artifacts—research summaries, structured plans—that are themselves reusable assets. Software engineer Dexter Horthy’s Research-Plan-Implement (RPI) (which has now evolved to a more elaborate QRSPI methodology)¹²¹ represents efforts to codify workflow decomposition in software engineering. Anthropic’s guidance on long-horizon tasks identifies three complementary techniques: compaction (summarizing history to preserve context across sessions), structured note-taking (maintaining process logs outside the context window), and multi-agent delegation.¹²²

Taken together, orchestration appears to support cognitive efficacy by structuring how context flows through human-AI workflows to maximize output quality while managing computational cost and error propagation. The design choices involved—agent architecture, process decomposition, context window management—create natural checkpoints for human judgment, ensuring that practitioners retain strategic control over decisions about work content and verification.

¹¹⁹ Rodrigo Olivares and Jiri De Jonghe, “The Chief of Staff Agent,” Claude Cookbook, September 12, 2025. Azeem Azhar, “Behind the Scenes of My AI Agent,” Exponential View, February 27, 2026. Commentarii Roamani, “Commentarii Roamani: Roam Depot Gems: Chief of Staff,” Roam Research, n.d., accessed April 19, 2026. Amazon Web Services, “How We Built an AI Chief of Staff for Enterprise Sellers Using Amazon Quick,” AWS Builder Center, March 18, 2026.

¹²⁰ The [Virtual Lab](#) project from Stanford’s Zou Group demonstrates the principle: an LLM principal investigator guides specialist agents (chemist, biologist, computational scientist) alongside a dedicated Scientific Critic that challenges outputs and flags errors. The system designed novel SARS-CoV-2 nanobodies experimentally validated as effective binders—a result attributable to orchestrated complementary perspectives, not any single agent’s capability.

¹²¹ Alex Lavaee, “From RPI to QRSPI: Rebuilding the First Structured Workflow for Coding Agents,” Blog, n.d.

¹²² Prithvi Rajasekaran, Ethan Dixon, Carly Ryan, and Jeremy Hadfield, “Effective Context Engineering for AI Agents,” Anthropic, September 29, 2025.

Exploration: reinvesting in new frontiers of value creation

Execution of effective specification and orchestration within a user-controlled computing environment appears to support efficiency (time and cost savings) without compromising—and indeed likely strengthening the quality of existing workflows and outputs. This paper uses “exploration” to refer to a practice—observed in lead users and the authors’ own experimentation—of reinvesting these gains in efficiency and quality into development of new forms of value creation and capabilities for shared problem-solving.

Increasing interactions with novelty

At the micro level of individuals and teams, exploration means redeploying AI’s efficiency dividend inside the workflow itself. For example, this could mean allocating more time to the interaction layer of the context web (e.g., human-to-human meetings and exchanges, new or deeper research into ideas and analyses; see Figure 4 and Table 2 above) relative to time that users spend interacting with machines (e.g., toil or busy work).

In the authors’ own implementation, context-maxxing has freed practitioners to allocate a higher proportion of total work time to structured human-to-human interaction. This aligns with the author’s earlier experimental work on vibe teaming, where reorganizing knowledge work around human-to-human collaboration served as a source of novelty and quality control, owing to the way in which dynamic team deliberation can help surface tacit knowledge and more diverse perspectives that might otherwise remain latent.¹²³ In general, these findings suggest that teams that deploy AI to take over more of human-machine “grunt work” (recording meetings, generating first drafts) are potentially able to devote a higher ratio of total work time to interacting with the novelty (human-to-human meetings and convenings, deep solo work, or structured dialogue with domain experts).

Expanding capabilities and infrastructures

Exploration may also entail moving from single-mode knowledge production to multi-stage or parallel capability and infrastructure development. This might include using AI to structure more sophisticated simulations or facilitations, building new analytical tools, or experimenting with forms of engagement that would not have been feasible under traditional resource constraints. The early signal is already visible in the one- and two-person companies reaching revenue scales that, a

¹²³ Jacob Taylor and Kershin Krishna, “Vibe teaming: How human-human-AI collaboration could disrupt knowledge work for the world’s toughest challenges,” *Global Economy and Development, Working Paper 193* (Washington, D.C.: Brookings Institution, June 2025).

decade ago, required 50 or 500 employees.¹²⁴ Meanwhile, Exponential View, a prominent emerging technology research studio, has reportedly developed a suite of new analytical dashboards and analyses without needing to increase headcount.^{125,126}

A software team's exploration dividend might fund new internal tooling—the tests, staging environments, and observability stacks previously too costly to maintain. A trades team's efficiency dividend could be reinvested in training the next apprentice, bidding on more sophisticated jobs, or increasing independence.¹²⁷

In sum, exploration supports mastery by directing efficiency gains toward activities that deepen human capability rather than simply accelerating existing production. When practitioners reinvest saved time into human-to-human interaction, deep research, or experimentation with new problem spaces, the resulting insights feed back into specification assets and orchestration patterns—driving the virtuous cycle described below. Exploration also supports collective efficacy: By creating space for the kinds of divergent, interpersonal exchange that surface tacit knowledge and diverse perspectives, it strengthens the conditions under which teams can solve problems that exceed any single user's capacity.

2.3 A potentially virtuous cycle for cognitive agency with AI?

Signs are emerging that these competences are mutually reinforcing. Specification produces reusable context artifacts that make orchestration more efficient. Effective orchestration generates novel insights that refine specification while freeing up human capacity for more interactions that lead to exploration of new ideas and opportunities. Exploration produces new expertise and specification assets that expand capabilities into new frontiers of value creation.

The virtuous cycle is how context-maxxing could generate compounding returns and cognitive agency for users of generative AI. Context-maxxing produces externalized, durable assets that accumulate and potentially appreciate through use. Unlike use of proprietary AI platform interfaces, context-maxxing builds persistent user-controlled digital infrastructure that reduces the marginal cost of future knowledge production while also requiring a new set of competencies for exercising control, efficacy, and mastery with generative AI.

Part 3 considers the implications of this framework for public policy.

¹²⁴ Erin Griffith, "How A.I. Helped One Man (and His Brother) Build a \$1.8 Billion Company," *The New York Times*, April 2, 2026.

¹²⁵ Azeem Azhar and Marija Gavrillov, "Seven Lessons from Building with AI," Exponential View, April 16, 2025, accessed April 18, 2026, <https://www.exponentialview.co/p/seven-lessons-from-automating-our>.

¹²⁶ Azeem Azhar, "The \$100 Trillion Productivity Puzzle," Exponential View, June 19, 2025, accessed April 18, 2026, <https://www.exponentialview.co/p/ai-is-ready-is-your-company>.

¹²⁷ Daron Acemoglu, David Autor, and Simon Johnson, *Building pro-worker artificial intelligence* (Washington, DC: The Hamilton Project at Brookings, February 2026).

Part 3. Discussion

This paper analyzed how open-source, user-controlled AI deployments could support cognitive agency, defined as the capacity of individuals and collectives to think and act with generative AI in ways that support control (shaping the informational environment or context), efficacy (value creation and collaborating effectively for shared problem-solving), and mastery (accumulating durable capability over time). By contrast, current evidence suggests that proprietary deployment architectures are associated with cognitive erosion, potentially linked to the way in which these systems concentrate context control in vendor hands.

The viral success of OpenClaw in early 2026 demonstrated that a fundamentally different deployment paradigm—one placing context control in user hands—is now feasible and accessible. Analysis of lead user implementations suggests that AI's deployment architecture should be treated as a critical variable in determining whether AI expands or narrows cognitive agency. For this reason, how AI is designed and deployed deserves to be a more prominent topic of public scrutiny and debate.¹²⁸

The playbook presented in Part 2 identifies five infrastructure building blocks and three reinforcing competencies that, taken together, constitute an emerging user-controlled computing environment for generative AI. Drawing on the analytical framework developed in Part 1, this paper proposes that context-maxxing could support cognitive agency across all three dimensions—control, efficacy, and mastery (see Table 4). Testing these predictions requires stronger evidence than the observational and self-reported accounts that currently shape much of the discussion around user-controlled, open-source AI deployments. At a minimum, three lines of empirical inquiry are needed.

1. **Matched-task comparisons across deployment architectures.** Different AI deployment environments (proprietary single-chat interfaces like ChatGPT, closed-agent systems like Claude Code or Replit, and user-controlled context-maxxing architectures) could be used to execute the same knowledge work tasks—such as policy memo drafting, research synthesis, or coding.
2. **Longitudinal tracking of capability development.** Some predictions associated with cognitive control, efficacy, and mastery need to be evaluated over time. Future research should examine whether context assets specification assets genuinely appreciate through use, whether orchestration patterns stabilize and improve, and whether users develop transferable competencies.

¹²⁸ Jacob Taylor, "The most important question when designing AI," Brookings Institution, May 20, 2024, <https://www.brookings.edu/articles/the-most-important-question-when-designing-ai/>.

3. **Institutional and collective-level analysis.** The playbook described here focuses primarily on individual and small-team implementation. Whether context-maxxing supports cognitive agency at organizational or societal scales—including effects on collective intelligence, cognitive diversity, and democratic resilience—remains speculative. Research at these scales would need to examine how context webs interact across teams, how firms reorganize to accommodate user-controlled computing environments for generative AI, and whether the infrastructure and competencies described here are accessible to practitioners beyond the technically proficient early adopters who currently dominate the community.

Table 4. Proposed mechanisms linking context-maxxing to cognitive agency

| Dimension | Expected effect on cognitive agency | Mechanism |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Control | Users will report and demonstrate greater perceived and actual control over the informational environment of their AI interactions compared to users of default proprietary interfaces. | Inspectable configuration files, multiple LLM providers, local data storage, and portability shift control over context assembly from vendor to user. |
| Efficacy | Users will produce equivalent or higher-quality outputs with less time devoted to establishing context and more time devoted to interaction with novelty (e.g., human-to-human interaction or deep work). | Reusable specification assets and structured orchestration reduce redundant context assembly; process decomposition creates efficiencies that can be reallocated to interactions with novelty that surface new insights, tacit knowledge, and diverse perspectives. |
| Mastery | Users will accumulate more durable and portable context assets and competencies that increase in value through repeated use across tasks and settings. | Specification externalizes domain knowledge into reusable templates; users take responsibility for optimizing orchestration patterns; context webs accumulate knowledge that compounds across applications. |

Growing new competencies for pro-human AI

Cognitive agency with generative AI suggests not just preserving autonomy over existing human capabilities, but developing entirely new capacities for operating as human-AI hybrid systems.¹²⁹ Millions of formerly non-technical people globally are now using open-source harnesses to make domain knowledge machine-legible (specification), route context through multi-agent architectures (orchestration), and—when it all works—reinvest efficiency gains from these processes into work on new problem spaces (exploration). These competencies did not exist beyond software development three years ago; no established pedagogy yet exists for teaching them at scale or measuring the effects of their acquisition. User-controlled AI deployment environments may

¹²⁹ Guszczka, et. al. Hybrid Intelligence: A Paradigm for More Responsible Practice (October 12, 2022). Available at SSRN: <https://ssrn.com/abstract=4301478>.

therefore be an important testbed for surfacing and cultivating novel competencies for human-AI collaboration.

This argument is relevant to recent suggestions that verification—the capacity to evaluate AI-generated outputs for accuracy, appropriateness, and alignment with intent—may become a defining competency for value capture through generative AI. For instance, Catalini, Hui, and Wu draw attention to a structural asymmetry between rapidly falling costs of AI-driven execution of work and increasing costs of humans verifying that work due to biological constraints of human time and attention.¹³⁰ Their model suggests a measurability gap where AI agents can execute faster than humans can meaningfully assess quality. In domains where AI outputs cannot be mechanically verified—i.e., judgment-dependent knowledge work—economic value accrues primarily to actors who bear the risk of verification: those willing to stake reputation, liability, or resources on the correctness of AI-assisted work.

Under default assumptions common to most recent economic modeling on this topic, that AI deployment infrastructure and human cognitive capacity are understood as fixed constraints,¹³¹ verification capacity has been shown to concentrate centrally in firms rather than workers.¹³² Firms aggregate context, fine-tune models, and build institutional quality-control systems—creating economies of scale in verification that individual workers cannot match.

This paper does not dispute the worrying trend implied by these analyses. But it does question the assumptions that shape them. AI deployment infrastructure is not neutral—it is a design variable that can determine whether AI expands or erodes human agency.¹³³ Likewise, and as research spanning human-AI interaction, computer-supported cooperative work, and collective intelligence science demonstrates, the biological limits to human verification are far from fixed by the computational capacity or bounded

¹³⁰ Christian Catalini, Xiang Hui, and Jane Wu, “Some Simple Economics of AGI,” arXiv preprint, February 25, 2026, <https://arxiv.org/abs/2602.20946>.

¹³¹ For example, Catalini, Hui and Wu (2026) frame the transition as the collision of two racing cost curves—an exponentially decaying cost to automate set against a biologically bottlenecked cost to verify. Brynjolfsson and Hitzig (2025) build the same comparative logic into their centralization argument based on individual bounded-rationality. Acemoglu et al. (2026) cast AI capability as rising along a trajectory set by frontier firms while worker capacity is either held static or granted modest “expertise-levelling.”

¹³² Erik Brynjolfsson and Zoë Hitzig, “AI’s Use of Knowledge in Society,” in *The Economics of Transformative AI*, National Bureau of Economic Research, 2025.

¹³³ To be clear, labor economists distinguish between [automation \(task replacement\) and augmentation \(task complementation\)](#), and recent work explicitly models “[task reinstatement](#)”—the creation of new tasks where humans retain comparative advantage. But these frameworks largely abstract from the architectural details determining whether a given tool cultivates user agency or erodes it. The question of how complementation happens—whether through proprietary black-box systems or user-controlled infrastructure that promotes agency—does [not typically arise as a design variable within these models](#). This importance of this question is visible in the [case of telephone switchboard operators](#): early expansion of the Bell network created ever more operator jobs, but operators’ own skill in routing calls, troubleshooting failures, and training customers made the system reliable and ubiquitous enough that automated switching became profitable to deploy, at which point the same infrastructure that had amplified their labor rapidly rendered the occupation nearly obsolete.

rationality of an individual brain; rather they can be systematically augmented through development of (shared) infrastructure, institutions, and capabilities that support synergies between humans and technologies.¹³⁴

Human-led verification of AI outputs is more credible, reliable, and reproducible where users can own, inspect, and govern the workflows that produce AI outputs, rather than relying on opaque systems they cannot meaningfully examine.¹³⁵ By placing context more meaningfully under user control, context-maxxing potentially establishes the conditions under which humans can expand verification capacity.

Through encouraging greater user control and responsibility over the deployment environment of generative AI, context-maxxing supports the creation of reusable specification assets and orchestration patterns that potentially appreciate with use. This is not wholly unique to context-maxxing: Disciplined users of proprietary systems can also refine context files, prompts, and project configurations. The distinctive features of context-maxxing are that these assets are maintained outside a single vendor interface, making them more editable and portable across models, and reusable across workflows. These features could reduce the marginal costs of future knowledge production while maintaining (and arguably strengthening) user capacity to inspect and verify the quality of AI outputs generated within such an environment. This raises the prospect that rather than ceding verification to platforms or firms, people accumulate portable expertise that increases their capacity to verify.

More research into the relationship between context-maxxing, cognitive agency, and verification capability could help address a tension in debates about the economics of AI. For instance, few economists would disagree with the need for what Acemoglu and colleagues refer to as pro-worker AI.¹³⁶ Yet, as some observers rightly point out, it is currently difficult to clearly define what “pro-worker AI” actually means in practice.¹³⁷ Context-maxxing would suggest that pro-worker AI is AI that is designed and deployed to expand human control, efficacy, and mastery in work, rather than extract user context or erode worker agency. Lead-users are helping demonstrate what the infrastructure building blocks and competencies of pro-worker—or pro-human AI more generally—

¹³⁴ James A. Evans, Benjamin H. Bratton, and Blaise Agüera y Arcas, “Agentic AI and the Next Intelligence Explosion,” *Science* 391, no. 6791 (March 19, 2026): eaeg1895, <https://www.science.org/doi/10.1126/science.aeg1895>. Indeed, for biological life, [all intelligence is collective intelligence](#). Humans appear individually bounded but are socially wired—producing [rational intelligence at the collective level](#). Cultural evolution research shows human success lies not in individual cognitive capacity but in “[collective brains](#)”: groups that interconnect, learn from one another, and accumulate knowledge exceeding what any individual could produce alone.

¹³⁵ This workflow-level verification complements, but cannot replace, evaluation and verification of the underlying LLMs themselves.

¹³⁶ Daron Acemoglu, David Autor, and Simon Johnson, “Building Pro-Worker AI,” Brookings Institution, February 23, 2026, <https://www.brookings.edu/articles/building-pro-worker-ai/>.

¹³⁷ Joshua Gans, “WTF Is Pro-Worker AI,” Joshua Gans’s Newsletter (Substack), February 26, 2026.

could look like if open-source, user-controlled AI deployment environments were to proliferate across teams, firms, and societies.

How to ‘commoditize’ user-controlled AI deployment infrastructure

At the time of writing, the infrastructure requirements and setup and running costs associated with context-maxxing are not trivial. Meanwhile, proprietary systems are continuing to improve their functionality and convenience, making it difficult for bottom-up movements alone to counter these network effects and economies of scale. Without deliberate (public) support to make user-controlled deployments more accessible to more people, these approaches risk remaining confined to users with existing technical skills or resources, while those less well-resourced remain takers of proprietary deployments that foreclose on context control—either by design or by default.

To this end, policy responses to OpenClaw in China—including what are termed “lobster policies” in Shenzhen and Wuxi offering up to 5 million yuan in compute subsidies—signal government willingness to treat open harnesses as a form of subsidized public infrastructure. Indeed, prior to the OpenClaw movement, public AI infrastructure efforts were already underway at different levels, including targeting public provisioning of compute, models, and datasets. The Public AI inference utility, for example, currently hosts open-source public interest models for free, powered by federated compute from donations from over six countries. Preliminary work has also established principles of public AI data commons and Public AI data flywheels that would allow users to benefit collectively from the value their data generates rather than ceding it to platforms.^{138,139} These initiatives point toward the possibility of a coherent agenda where investments in federated compute, privacy-preserving machine learning architectures, and horizontal infrastructure (portability standards, shared context webs, public-interest model tiers) could combine to scale user-controlled systems and community-driven applications.¹⁴⁰

As David Eaves has argued in the context of promoting agency in society-scale digital infrastructure like cloud computing, if cognitive agency with generative AI is the aim, the priority should be “commoditizing” the AI deployment infra to make it as simple as possible for users to exercise control over their context.¹⁴¹

¹³⁸ Mozilla Foundation. “Mozilla Data Collective Redefines How AI Data Is Created, Shared, and Who Benefits.” Press release, November 6, 2025. <https://www.mozilla.org/en/meet-mozilla/press-center/mozilla-data-collective-launches/>.

¹³⁹ Vincent, Nick. Public AI Data Flywheel: An Open Mini-Book and Prototype Implementation. GitHub Pages, 2025. https://nickmvincent.github.io/paidf_consultation/

¹⁴⁰ Taylor, Jacob, Thomas Kehler, Alex “Sandy” Pentland, and Martin Reeves. “Three Principles for Growing an AI Ecosystem That Works for People and Planet.” Brookings Institution, August 1, 2025. <https://www.brookings.edu/articles/three-principles-for-growing-an-ai-ecosystem-that-works-for-people-and-planet/>.

¹⁴¹ Eaves, David. “The Path to a Sovereign Tech Stack Is via a Commodified Tech Stack.” Tech Policy Press, December 15, 2025. <https://techpolicy.press/the-path-to-a-sovereign-tech-stack-is-via-a-commodified-tech-stack/>.

One key area determining user control is the interoperability between protocols and standards across agent harnesses. The good news is that Model Context Protocol (MCP; now adopted by Microsoft, OpenAI, Anthropic, Gemini, Baidu, and Alibaba) appears to be emerging as a standardized layer for tool calls and memory—like the “USB-C port for AI applications.”^{142,143} The Human Context Protocol proposes extending MCP to standardize formats for user-controlled context portability across agent platforms. The Loyal Agents initiative similarly explores protocols and guardrails that would mean AI agents maintain fidelity to user interests rather than platform incentives.¹⁴⁴ Together, these protocols could provide a common layer for user-controlled memory, tool calls, and governance that can sit underneath many different harnesses and model vendors.

Top-down signals of incentives and investment will be important to create the conditions that allow user-focused protocols to survive and scale. This will likely require treating user-controlled deployment infrastructure as a new form of (digital) public good or infrastructure worthy of investment—particularly for sustainable development contexts where cost and capacity are often the binding constraint.¹⁴⁵ Decisionmakers can learn many lessons from longstanding efforts to grow digital public goods and digital public infrastructure (DPI) as a set of predominantly open-source, interoperable, and publicly governed applications and protocols for supporting inclusive public service delivery and equitable digital economic activity.¹⁴⁶ India’s emerging “AI-DPI” agenda shows what it looks like to apply DPI thinking directly to AI.¹⁴⁷ Under the IndiaAI mission, initiatives such as IndiaAI Kosh and the IndiaAI Compute Portal treat models, domain datasets, and subsidized GPU clusters as digital public goods, while projects like Telangana’s TGDEx and the proposed “Unified AI Layer” (“UPI for AI”) build horizontal data and API rails that sector-specific teams can plug into for education, health, agriculture, and other use cases.¹⁴⁸

¹⁴² Girija, Sanjay Surendranath, and Lakshit Arora. “MCP: The Universal Connector for Building Smarter, Modular AI Agents.” InfoQ, August 28, 2025. <https://www.infoq.com/articles/mcp-connector-for-building-smarter-modular-ai-agents/>.

¹⁴³ Das, Shanaya. “China’s MCP Adoption: The Rise of AI Assistants That Actually Do Things.” AI News, April 26, 2025. <https://www.artificialintelligence-news.com/news/chinas-mcp-adoption-ai-assistants-that-actually-do-things/>.

¹⁴⁴ Stanford Digital Economy Lab and Consumer Reports Innovation Lab. Authentic AI: Bringing Authentication, Authorization, and Identity into the AI Agent World. Loyal Agents white paper, 2026. <https://hcp.loyalagents.org/authentic-AI-whitepaper.pdf>.

¹⁴⁵ Chia, Han Sheng. “Cutting Through the Noise: Reimagining Tech for Good.” Blog post, Center for Global Development, March 15, 2026. <https://www.cgdev.org/blog/cutting-through-noise-reimagining-tech-good>.

¹⁴⁶ Eaves, David, and Jordan Sandman. “What Is Digital Public Infrastructure?” Co-Develop, October 20, 2023. <https://www.codevelop.fund/insights-1/what-is-digital-public-infrastructure>

¹⁴⁷ Kapoor, Shalini, and Tanvi Lall. “The Possibilities of DPI & AI.” People+AI Blog, November 5, 2025. <https://peopleplus.ai/blog/the-possibilities-of-dpi-ai>.

¹⁴⁸ Economic Times, “Explained: IndiaAI Compute Portal, AIKosha and Other Initiatives Under the IndiaAI Mission,” Economic Times, March 6, 2025, https://m.economictimes.com/tech/technology/explained-indiaai-compute-portal-aikosha-and-other-initiatives-under-the-indiaai-mission/amp_articles/118780355.cms; Swati Bharadwaj, “Telangana Rolls Out TGDEx, India’s First State-Led Digital Public Infrastructure for AI,”

From context-maxxing to context gardening

This paper adopted the language of context-maxxing deliberately—to meet an emerging industry discourse on its own terms and to make legible what a growing community of practitioners is already building. But the framework described here ultimately points beyond the logic of optimization that the "-maxxing" suffix implies, to a new kind of relationship between people and AI systems that instead requires organic growth—and surrounding infrastructure designed to support that growth. In this spirit, a truer name for the approaches described in this paper might be something closer to “context gardening”—a practice oriented not toward algorithmic extraction or optimization, but toward the patient cultivation of novel human capabilities with AI.

Times of India, July 2, 2025, <https://timesofindia.indiatimes.com/business/india-business/telangana-rolls-out-tgdex-indias-first-state-led-digital-public-infrastructure-for-ai/articleshow/122210745.cms>; https://www.oecd.org/en/publications/digital-public-infrastructure-for-digital-governments_ff525dc8-en.html; OpenAgriNet. “OpenAgriNet (OAN): Building Digital Agriculture Grids Across the Globe.” OpenAgriNet, 2024. <https://openagrinet.global>.

References

Acemoglu, Daron, David Autor, and Simon Johnson. Building pro-worker artificial intelligence. Washington, DC: The Hamilton Project at Brookings, February 2026.

Acemoglu, Daron, David Autor, and Simon Johnson. "Building Pro-Worker AI." Brookings Institution, February 23, 2026. <https://www.brookings.edu/articles/building-pro-worker-ai/>.

Acemoglu, Daron, Tianyi Lin, Asuman Ozdaglar, and James Siderius. How AI Aggregation Affects Knowledge. NBER Working Paper No. 35036. Cambridge, MA: National Bureau of Economic Research, April 2026. <https://www.nber.org/papers/w35036>.

Ali, Mahadi Mokbul. "Cognitive Sovereignty: A Theory and Initial Validation of Human Autonomy in Algorithmic Decision Systems." Research Square, version 1, posted March 17, 2026. <https://www.researchsquare.com/article/rs-9145237/v1>.

Amazon Web Services. "Edge AI and Global Inference Distribution." AWS Prescriptive Guidance. Accessed April 15, 2026. <https://docs.aws.amazon.com/prescriptive-guidance/latest/agent-ai-serverless/edge-ai.html>.

Amazon Web Services. "How We Built an AI Chief of Staff for Enterprise Sellers Using Amazon Quick." AWS Builder Center, March 18, 2026.

Andrea. "Anthropic Shifts API Strategy for OpenClaw, Ending Pro Subscription Access for Third-Party Agents." Hyperright, April 6, 2026.

Andrea. "OpenClaw Smashes Records: The Viral AI Agent Is Shaking Up GitHub." Hyperright, January 30, 2026. <https://hyperright.com/openclaw-ai-assistant-rebrand/>.

Apple. "Buy Mac mini." Apple Store. Accessed April 15, 2026. <https://www.apple.com/shop/buy-mac/mac-mini>.

Azhar, Azeem. "Behind the Scenes of My AI Agent." Exponential View, February 27, 2026. <https://www.exponentialview.co/p/behind-the-scenes-of-my-ai-agent>.

Azhar, Azeem. "Seven Lessons from Building with AI." Exponential View, April 16, 2025. Accessed April 18, 2026. <https://www.exponentialview.co/p/seven-lessons-from-automating-our>.

Azhar, Azeem. "The \$100 Trillion Productivity Puzzle." Exponential View, June 19, 2025. Accessed April 18, 2026. <https://www.exponentialview.co/p/ai-is-ready-is-your-company>.

Azhar, Azeem. "Why I Changed My Mind about Apple." Exponential View, March 19, 2026. <https://www.exponentialview.co/p/why-i-changed-my-mind-about-apple>.

Banh, L., and G. Strobel. "Generative Artificial Intelligence." Electronic Markets 33, no. 63 (2023). <https://doi.org/10.1007/s12525-023-00680-1>.

Bandura, Albert. "Social Cognitive Theory: An Agentic Perspective." Annual Review of Psychology 52 (2001): 1–26.

Barcaui, André. "ChatGPT as a Cognitive Crutch: Evidence from a Randomized Controlled Trial on Knowledge Retention." Social Sciences & Humanities Open 12 (2025): 102287. <https://www.sciencedirect.com/science/article/pii/S2590291125010186>.

Bergmann, Dave. "What Is a Context Window?" IBM. Accessed April 20, 2026. <https://www.ibm.com/think/topics/context-window>.

BetterUp Labs and Stanford Social Media Lab. "Workslop: The Hidden Cost of AI-Generated Busywork." BetterUp Labs, September 20, 2025. <https://www.betterup.com/workslop>.

Bharadwaj, Swati. "Telangana Rolls Out TGDex, India's First State-Led Digital Public Infrastructure for AI." Times of India, July 2, 2025. <https://timesofindia.indiatimes.com/business/india-business/telangana-rolls-out-tgdex-indias-first-state-led-digital-public-infrastructure-for-ai/articleshow/122210745.cms>.

Bousquette, Isabelle. "Why Some Companies Say AI 'Tokenmaxxing' Is Key to Survival." Wall Street Journal, April 14, 2026.

Bracken, Mike, David Eaves, and Michelle Wronski. The Service Gap: Europe's International Digital Strategy 2025. Brussels: The Lisbon Council, 2025. <https://lisboncouncil.net/the-service-gap-europe-international-digital-strategy-2025/>.

Brainard, Jeffrey. "AI Algorithms Can Become 'Agents of Chaos.'" Science, March 23, 2026. <https://www.science.org/content/article/ai-algorithms-can-become-agents-chaos>.

Brynjolfsson, Erik, and Zoë Hitzig. "AI's Use of Knowledge in Society." In The Economics of Transformative AI. National Bureau of Economic Research, 2025.

Cang, Wei. "Wuxi High-Tech Zone Launches Supportive Policies for OpenClaw Project." China Daily, March 10, 2026. <https://www.chinadaily.com.cn/a/202603/10/WS69afd0f6a310d6866eb3d021.html>.

Catalini, Christian, Xiang Hui, and Jane Wu. "Some Simple Economics of AGI." arXiv preprint, February 25, 2026. <https://arxiv.org/abs/2602.20946>.

Chaudhary, Yaqub, and Jonnie Penn. "Beware the Intention Economy: Collection and Commodification of Intent via Large Language Models." *Harvard Data Science Review* 6, no. 4 (December 30, 2024). <https://hdsr.mitpress.mit.edu/pub/ujvharkk>.

Chayka, Kyle. "Why Tech Bros Are Now Obsessed with Taste." *The New Yorker*, March 18, 2026. <https://www.newyorker.com/culture/infinite-scroll/why-tech-bros-are-now-obsessed-with-taste>.

Chen, Yi, Fei Li, and Marcel Preuss. "Algorithmic Attention and Content Creation on Social Media Platforms." SSRN, March 17, 2025. <https://ssrn.com/abstract=5182754>.

Chia, Han Sheng. "Cutting Through the Noise: Reimagining Tech for Good." Center for Global Development, March 15, 2026. <https://www.cgdev.org/blog/cutting-through-noise-reimagining-tech-good>.

Chirayath, Ginto, K. Premamalini, and Jeena Joseph. "Cognitive Offloading or Cognitive Overload? How AI Alters the Mental Architecture of Coping." *Frontiers in Psychology* 16 (2025): 1699320. <https://www.frontiersin.org/journals/psychology/articles/10.3389/fpsyg.2025.1699320/full>.

Ciriello, Raffaele F., and Kathryn Backholer. "OpenAI Will Put Ads in ChatGPT. This Opens a New Door for Dangerous Influence." *The Conversation*, January 23, 2026. <https://doi.org/10.64628/AA.e3pvkfugk>.

Claburn, Thomas. "Anthropic Tweaks Timed Usage Limits to Discourage Claude Demand during Peak Hours." *The Register*, March 26, 2026. https://www.theregister.com/2026/03/26/anthropic_tweaks_usage_limits/.

Commentarii Roamani. "Commentarii Roamani: Roam Depot Gems: Chief of Staff." Roam Research, n.d. Accessed April 19, 2026.

Csikszentmihalyi, Mihaly. *Flow: The Psychology of Optimal Experience*. New York: Harper & Row, 1990.

Cutler, Silas. "OpenClaw in the Wild: Mapping the Public Exposure of a Viral AI Assistant." Censys, January 31, 2026. <https://censys.com/blog/openclaw-in-the-wild-mapping-the-public-exposure-of-a-viral-ai-assistant/>.

Danaher, John, et al. "Algorithmic Governance: Developing a Research Agenda through the Power of Collective Intelligence." *Big Data & Society* 4, no. 2 (2017). <https://doi.org/10.1177/2053951717726554>.

Das, Shanaya. "China's MCP Adoption: The Rise of AI Assistants That Actually Do Things." AI News, April 26, 2025. <https://www.artificialintelligence-news.com/news/chinas-mcp-adoption-ai-assistants-that-actually-do-things/>.

Deci, Edward L., and Richard M. Ryan. "The 'What' and 'Why' of Goal Pursuits: Human Needs and the Self-Determination of Behavior." *Psychological Inquiry* 11, no. 4 (2000): 227–68.

Dell'Acqua, Fabrizio, Charles Ayoubi, Hila Lifshitz-Assaf, Raffaella Sadun, Ethan R. Mollick, Lilach Mollick, Yi Han, Jeff Goldman, Hari Nair, Stew Taub, and Karim R. Lakhani. "The Cybernetic Teammate: A Field Experiment on Generative AI Reshaping Teamwork and Expertise." Harvard Business School Strategy Unit Working Paper No. 25-043, March 28, 2025. SSRN. <https://ssrn.com/abstract=5188231>.

Doleac, Jennifer, Anna Harvey, Jonathon Attridge, Erin Dalton, Dan Kreisman, Weston Merrick, Anthony F. Pipa, Jenni Owen, and Jim Sullivan. "Building State and Local Government Innovation Capacity: Investing in University–Government Innovation Partnerships." 17 Rooms, Brookings Institution and Rockefeller Foundation, May 16, 2025. <https://www.brookings.edu/articles/building-state-and-local-government-innovation-capacity-investing-in-university-government-innovation-partnerships/>.

Dubois, Patric. "The Context Flywheel: Why the Best AI Coding Teams Will Win on Context." Tessel Blog. Accessed April 15, 2026. <https://tessl.io/blog/the-context-flywheel-why-the-best-ai-coding-teams-will-win-on-context/>.

Eaves, David. "The Path to a Sovereign Tech Stack Is via a Commodified Tech Stack." Tech Policy Press, December 15, 2025. <https://techpolicy.press/the-path-to-a-sovereign-tech-stack-is-via-a-commodified-tech-stack/>.

Eaves, David, and Jordan Sandman. "What Is Digital Public Infrastructure?" Co-Develop, October 20, 2023. <https://www.codevelop.fund/insights-1/what-is-digital-public-infrastructure>.

Economic Times. "Explained: IndiaAI Compute Portal, AIKosha and Other Initiatives Under the IndiaAI Mission." *Economic Times*, March 6, 2025. https://m.economictimes.com/tech/technology/explained-indiaai-compute-portal-aikosha-and-other-initiatives-under-the-indiaai-mission/amp_articleshow/118780355.cms.

Ericsson, K. Anders, Ralf Th. Krampe, and Clemens Tesch-Römer. "The Role of Deliberate Practice in the Acquisition of Expert Performance." *Psychological Review* 100, no. 3 (1993): 363–406.

Evans, James A., Benjamin H. Bratton, and Blaise Agüera y Arcas. "Agentic AI and the Next Intelligence Explosion." *Science* 391, no. 6791 (March 19, 2026): eaeg1895.

<https://www.science.org/doi/10.1126/science.aeg1895>.

Fetterman, Adam G., John B. Wilkerson, and Benjamin J. Blankenship. "From Offloading to Engagement: An Experimental Study on Structured Prompting and Critical Reasoning with Generative AI." *Data* 10, no. 11 (2025): 172.

Flavell, John H. "Metacognition and Cognitive Monitoring: A New Area of Cognitive-Developmental Inquiry." *American Psychologist* 34, no. 10 (1979): 906–11.

Floridi, Luciano. *The Ethics of Information*. Oxford: Oxford University Press, 2013.

Fosters, Tom. "Don't Get Pinched: The OpenClaw Vulnerabilities." Kaspersky Official Blog, February 10, 2026. <https://www.kaspersky.com/blog/openclaw-vulnerabilities-exposed/55263>.

FounderCoHo. "OpenClaw: The Malware You Installed on Purpose — A Security Playbook: A Layer-by-Layer Risk Analysis of the Fastest-Growing Open-Source AI Agent in History — For Users, Builders, and Maintainers." FounderCoHo, March 14, 2026.

<https://foundercoho.substack.com/p/openclaw-the-malware-you-installed>.

Gans, Joshua. "WTF Is Pro-Worker AI." Joshua Gans's Newsletter, February 26, 2026.

Gartner. "Context Engineering." Accessed April 15, 2026.

<https://www.gartner.com/en/articles/context-engineering>.

George, Rachel, and Ian Klaus. *AI and Democracy: Mapping the Intersections*.

Washington, DC: Carnegie Endowment for International Peace, 2026.

<https://carnegieendowment.org/research/2026/01/ai-and-democracy-mapping-the-intersections>.

Georgiou, Georgios P. "ChatGPT Produces More 'Lazy' Thinkers: Evidence of Cognitive Engagement Decline." *arXiv* preprint arXiv:2507.00181, June 30, 2025.

<https://doi.org/10.48550/arXiv.2507.00181>.

GetDeploying. "VPS Price Comparison." Updated March 26, 2026.

<https://getdeploying.com/reference/compute-prices>.

Girija, Sanjay Surendranath, and Lakshit Arora. "MCP: The Universal Connector for Building Smarter, Modular AI Agents." *InfoQ*, August 28, 2025.

<https://www.infoq.com/articles/mcp-connector-for-building-smarter-modular-ai-agents/>.

GitHub Advisory Database. "OpenClaw Affected by SSRF via Attachment/Media URL Hydration." GitHub, published February 17, 2026. <https://github.com/advisories/GHSA-wfp2-v9c7-fh79>.

GitHub Advisory Database. "OpenClaw Is Vulnerable to Path Traversal through Path Validation Bypass." GitHub, published March 26, 2026. <https://github.com/advisories/GHSA-hggm-x7r9-mm7v>.

Google Cloud. "Edge Hybrid Pattern." Cloud Architecture Center, last modified January 23, 2025. <https://docs.cloud.google.com/architecture/hybrid-multicloud-patterns-and-practices/edge-hybrid-pattern>.

Griffith, Erin. "How A.I. Helped One Man (and His Brother) Build a \$1.8 Billion Company." The New York Times, April 2, 2026.

Grinschgl, Sandra, Frank Papenmeier, and Hauke Meyerhoff. (2021). Consequences of cognitive offloading: boosting performance but diminishing memory. *Q. J. Exp. Psychol.* 74, 1477–1496. doi: 10.1177/17470218211008060.

Guszcza, James, David Danks, Craig R. Fox, Kristian J. Hammond, Daniel E. Ho, Alex Imas, James Landay, Margaret Levi, Jennifer Logg, Rosalind W. Picard, Manish Raghavan, Allison Stanger, Zachary Ugolnik, and Anita Williams Woolley. "Hybrid Intelligence: A Paradigm for More Responsible Practice." SSRN, October 12, 2022. <https://ssrn.com/abstract=4301478>.

Harré, Michael S., Catherine Drysdale, and Jaime Ruiz-Serra. "Theory of Mind Enhances Collective Intelligence." arXiv preprint arXiv:2411.09168, November 14, 2024. <https://arxiv.org/abs/2411.09168>.

HEmile. Obsidian Neo4j Graph View. GitHub repository. Accessed April 24, 2026. <https://github.com/HEmile/obsidian-neo4j-graph-view>.

Henrich, Joseph. *The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*. Princeton, NJ: Princeton University Press, 2016.

Herrman, John. "My Adventures With 'The AI That Actually Does Things.'" *New York Magazine*, April 28, 2026. <https://nymag.com/intelligencer/article/my-adventures-setting-up-openclaw-agent.html>.

Hong, Kelly, Anton Troynikov, and Jeff Huber. "Context Rot: How Increasing Input Tokens Impacts LLM Performance." Technical report, Chroma, July 2025. <https://www.trychroma.com/research/context-rot>.

Horthy, Dex. "Advanced Context Engineering for Coding Agents." HumanLayer Blog, August 29, 2025. Accessed April 19, 2026. <https://www.humanlayer.dev/blog/advanced-context-engineering>.

HumanLayer. "HumanLayer — Close Your Editor Forever." Accessed April 15, 2026. <https://www.humanlayer.dev/>.

Ide, Enrique, and Eduard Talamas. "Automation, Verification, and the Hollow Economy." Working paper, 2025.

Jones, Charles I., and Christopher Tonetti. "Nonrivalry and the Economics of Data." NBER Working Paper No. 26260, September 2019, revised April 2020. <https://www.nber.org/papers/w26260>.

Joren, Håkon Lønningdal, et al. "Sufficient Context: A New Lens on Retrieval-Augmented Generation Systems." In Proceedings of the 13th International Conference on Learning Representations, Vienna, 2025. <https://arxiv.org/abs/2411.06037>.

Kabeer, Naila. "Resources, Agency, Achievements: Reflections on the Measurement of Women's Empowerment." *Development and Change* 30, no. 3 (1999): 435–464. <https://doi.org/10.1111/1467-7660.00125>.

Kapoor, Shalini, and Tanvi Lall. "The Possibilities of DPI & AI." People+AI Blog, November 5, 2025. <https://peopleplus.ai/blog/the-possibilities-of-dpi-ai>.

Kaufmann, Rafael, Pranav Gupta, and Jacob Taylor. "An Active Inference Model of Collective Intelligence." *Entropy* 23, no. 7 (2021): 830. <https://doi.org/10.3390/e23070830>.

Kim, Jin, et al. "People Reduce Workers' Compensation for Using Artificial Intelligence (AI)." arXiv preprint arXiv:2501.13228, January 22, 2025. <https://arxiv.org/abs/2501.13228>.

Kim, Yubin, et al. "Towards a Science of Scaling Agent Systems." arXiv preprint, December 8, 2025. <https://arxiv.org/abs/2512.08296>.

Kosmyna, Nataliya, et al. "Your Brain on ChatGPT: Accumulation of Cognitive Debt When Using an AI Assistant for Essay Writing Task." MIT Media Lab, June 7, 2025. <https://arxiv.org/pdf/2506.08872.pdf>.

Kotler, Steven, Michael Mannino, Scott Kelso, and Richard Huskey. "First Few Seconds for Flow: A Comprehensive Proposal of the Neurobiology and Neurodynamics of State Onset." *Neuroscience & Biobehavioral Reviews* 143 (2022): 104956. <https://doi.org/10.1016/j.neubiorev.2022.104956>.

Krishnan, Rohit. "Why Coase Needs Hayek: Sometimes Smart Planners Lose to Simple Markets." *Strange Loop Canon*, May 2, 2026. <https://www.strangeloopcanon.com/p/why-smart-planners-lose-to-simple>.

Lavaee, Alex. "From RPI to QRSPI: Rebuilding the First Structured Workflow for Coding Agents." Blog, n.d.

Levy, Mosh, Alon Jacoby, and Yoav Goldberg. "Same Task, More Tokens: The Impact of Input Length on the Reasoning Performance of Large Language Models." In Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics, 2024. <https://arxiv.org/abs/2402.14848>.

Liu, Grace, Brian Christian, Tsvetomira Dumbalska, Michiel A. Bakker, and Rachit Dubey. "AI Assistance Reduces Persistence and Hurts Independent Performance." *arXiv* preprint arXiv:2604.04721, revised April 7, 2026. <https://doi.org/10.48550/arXiv.2604.04721>.

Lynn, Barry C. "The Big Tech Extortion Racket: How Google, Amazon, and Facebook Control Our Lives." *Harper's Magazine*, September 2020.

Maas, Matthijs M. "Sociotechnical Change: AI as Regulatory Rationale and Target." In *Architectures of Global AI Governance: From Technological Change to Human Choice*, edited by Wulf A. Kaal and Francesco Paolo Patti, chapter 6. Oxford: Oxford University Press, 2025. <https://doi.org/10.1093/9780191988455.003.0006>.

Majic, Josipa. "VCs Say Context Graphs Might Be the Next Big Thing in AI." *Forbes*, April 3, 2026. <https://www.forbes.com/sites/josipamajic/2026/04/03/vcs-say-context-graphs-might-be-the-next-big-thing-in-ai/>.

Mollick, Ethan. "Latent Expertise: Everyone Is in R&D." *One Useful Thing*, June 20, 2024. <https://www.oneusefulthing.org/p/latent-expertise-everyone-is-in-r>.

Mollick, Ethan. "The Bitter Lesson versus The Garbage Can." *One Useful Thing*, July 28, 2025.

Mozilla Foundation. "Mozilla Data Collective Redefines How AI Data Is Created, Shared, and Who Benefits." Press release, November 6, 2025. <https://www.mozillafoundation.org/en/meet-mozilla/press-center/mozilla-data-collective-launches/>.

Muthukrishna, Michael, and Joseph Henrich. "Innovation in the Collective Brain." *Philosophical Transactions of the Royal Society B: Biological Sciences* 371, no. 1690 (2016): 20150192.

National Vulnerability Database. "CVE-2026-25253 Detail." National Institute of Standards and Technology, published February 1, 2026. <https://nvd.nist.gov/vuln/detail/CVE-2026-25253>.

Noveck, Beth Simone. Solving Public Problems: A Practical Guide to Fix Our Government and Change Our World. New Haven, CT: Yale University Press, 2021.

OpenAI. "How Your Data Is Used to Improve Model Performance." OpenAI, updated April 28, 2025. <https://openai.com/policies/how-your-data-is-used-to-improve-model-performance/>.

OpenClaw. "openclaw/openclaw." GitHub repository, created January 29, 2026. <https://github.com/openclaw/openclaw>.

OpenClaw VPS. "OpenClaw Statistics 2026: Growth, Users, Security, Data." OpenClaw VPS Blog, April 2, 2026. <https://openclawvps.io/blog/openclaw-statistics>.

Oreopoulos, Philip, Oliver Keyes-Krysakowski, and Deepak Agarwal. How In-School Supervised Ed-Tech Support Produces Massive Learning Gains: A Khan Academy Field Experiment in India. NBER Working Paper No. 34683. Cambridge, MA: National Bureau of Economic Research, January 2026. <https://www.nber.org/papers/w34683>.

Ostrom, Elinor. Governing the Commons: The Evolution of Institutions for Collective Action. Cambridge: Cambridge University Press, 1990.

Page, Scott E. The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies. Princeton, NJ: Princeton University Press, 2007.

Pahlka, Jennifer. Recoding America: Why Government Is Failing in the Digital Age and How We Can Do Better. New York: Metropolitan Books, 2023.

Pelenc, Jérôme, and Catherine Ballet. "Is Amartya Sen's Sustainable Freedom a Broader Vision of Sustainability?" *Ecological Economics* 105 (September 2014): 1–7. <https://doi.org/10.1016/j.ecolecon.2014.05.010>.

Price Per Token. "LLM API Pricing 2026: Compare 300+ AI Model Costs." Accessed April 11, 2026. <https://pricepertoken.com/>.

Qin, Ying, Robert W. Smith, Mathieu Stillman, and Daniel M. Romero. "Generative AI Enhances Individual Creativity but Reduces the Collective Diversity of Novel Content." *Science Advances* 10, no. 28 (July 12, 2024): eadn5290. <https://pmc.ncbi.nlm.nih.gov/articles/PMC11244532/>.

Radsch, Courtney C. "The Battle for Cognitive Liberty in the Age of Corporate AI." Tech Policy Press, January 6, 2026. <https://www.techpolicy.press/the-battle-for-cognitive-liberty-in-the-age-of-corporate-ai/>.

Rajasekaran, Prithvi, Ethan Dixon, Carly Ryan, and Jeremy Hadfield. "Effective Context Engineering for AI Agents." Anthropic, September 29, 2025.

Randazzo, Chiara, Federico Dell'Acqua, Ethan Mollick, François Cadelon, and Karim R. Lakhani. "Cyborgs, Centaurs and Self-Automators." Harvard Business School Working Paper No. 26-036, 2025.

Reuters. "Meta Acquires AI Agent Social Network Moltbook." Reuters, March 10, 2026.

Reuters. "Tencent Integrates WeChat with OpenClaw AI Agent amid China Tech Battle." Reuters, March 22, 2026, updated March 23, 2026.

<https://www.reuters.com/technology/tencent-integrates-wechat-with-openclaw-ai-agent-amid-china-tech-battle-2026-03-22/>.

Riedl, Christoph, Saiph Savage, and Josie Zvelebilova. "Cognitive Spillover in Human-AI Teams." ACM Transactions on Computer-Human Interaction, Just Accepted, April 2026.

<https://doi.org/10.1145/3805039>.

Riedl, Christoph, Young Ji Kim, Pranav Gupta, Thomas W. Malone, and Anita Williams Woolley. "Quantifying Collective Intelligence in Human Groups." Proceedings of the National Academy of Sciences of the United States of America 118, no. 21 (May 25, 2021): e2005737118. <https://doi.org/10.1073/pnas.2005737118>.

Sahoo, Pranab, et al. "A Systematic Survey of Prompt Engineering in Large Language Models: Techniques and Applications." arXiv preprint, 2024.

<https://arxiv.org/abs/2402.07927>.

Science. "AI Algorithms Can Become Agents of Chaos." Science.

<https://www.science.org/content/article/ai-algorithms-can-become-agents-chaos>.

Sen, Amartya. Commodities and Capabilities. Amsterdam: North-Holland, 1985.

Sen, Amartya. Development as Freedom. New York: Alfred A. Knopf, 1999.

ServiceTitan. "HVAC Diagnostic Chart: What You Need to Know." March 10, 2025.

<https://www.servicetitan.com/blog/hvac-diagnostic-chart>.

Shah, Anand V., et al. "Robust AI Personalization Controls: The Human Context Protocol." SSRN Scholarly Paper No. 5403981, September 7, 2025.

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5403981.

Singh, Anjali, Karan Taneja, Zhitong Guan, and Avijit Ghosh. "Protecting Human Cognition in the Age of AI." arXiv preprint, last revised September 29, 2025.

<https://arxiv.org/abs/2502.12447>.

Sourati, Zhivar, Ali S. Ziabari, and Mohammad Dehghani. "The Homogenizing Effect of Large Language Models on Human Expression and Thought." *Trends in Cognitive Sciences*, advance online publication, 2026. <https://doi.org/10.1016/j.tics.2026.01.003>.

Stanford Digital Economy Lab and Consumer Reports Innovation Lab. *Authentic AI: Bringing Authentication, Authorization, and Identity into the AI Agent World. Loyal Agents white paper*, 2026. <https://hcp.loyalagents.org/authentic-AI-whitepaper.pdf>.

Star History. "OpenClaw Star History." Accessed April 24, 2026. <https://www.star-history.com/openclaw/openclaw#history>.

Stockmeyer, Otto. "It's All About IRAC." *Cooley Law School*. Accessed April 30, 2026.

<https://cooley.edu/blog/its-all-about-irac>.

Swain, Gyana. "OpenAI Hires OpenClaw Founder as AI Agent Race Intensifies." *InfoWorld*, February 16, 2026.

Tanner, Matt. "Enterprise Software Architecture Patterns: The Complete Guide."

vFunction, June 18, 2025. <https://vfunction.com/blog/enterprise-software-architecture-patterns/>.

Taylor, Jacob, and Kershlin Krishna. *Vibe Teaming: How Human-Human-AI Collaboration Could Disrupt Knowledge Work for the World's Toughest Challenges*. Working Paper 193. Washington, DC: Brookings Institution, June 2025.

Taylor, Jacob, and Scott E. Page. "AI Is Changing the Physics of Collective Intelligence—How Do We Respond?" Brookings Institution, December 16, 2025.

<https://www.brookings.edu/articles/ai-is-changing-the-physics-of-collective-intelligence-how-do-we-respond/>.

Tirole, Jean. "Competition and the Industrial Challenge for the Digital Age." Background paper for the IFS Deaton Review, April 3, 2020. https://www.tse-fr.eu/sites/default/files/TSE/documents/doc/by/tirole/competition_and_the_industrial_challenge_april_3_2020.pdf.

TLDL. "OpenAI API Pricing 2026 — GPT-5.4, O3, O1 & GPT-4o Cost Per Token." Updated March 5, 2026. Accessed April 11, 2026. <https://www.tldl.io/resources/openai-api-pricing>.

United Nations Development Programme. Human Development Report 2025: A Matter of Choice — People and Possibilities in the Age of AI. New York: UNDP, 2025.

<https://hdr.undp.org/content/human-development-report-2025>.

Vaswani, Ashish, et al. "Attention Is All You Need." arXiv preprint, first posted June 12, 2017. <https://arxiv.org/abs/1706.03762>.

Vincent, Nick. Public AI Data Flywheel: An Open Mini-Book and Prototype Implementation. GitHub Pages, 2025. https://nickmvincent.github.io/paidf_consultation/.

Walker, Aidan. "Clavicular and Contentmaxxing: The Next Step After Groyperfication." How to Do Things With Memes, January 20, 2026. <https://howtodothingswithmemes.substack.com/p/clavicular-and-contentmaxxing>.

Westby, Samuel, and Christoph Riedl. "Collective Intelligence in Human-AI Teams: A Bayesian Theory of Mind Approach." In Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence, 2023. <https://arxiv.org/abs/2208.11660>.

Whitten, Allison. "New Tool Helps AI and Humans Learn To Code Better." Stanford Institute for Human-Centered Artificial Intelligence, May 1, 2023.

<https://hai.stanford.edu/news/new-tool-helps-ai-and-humans-learn-code-better?sf180948227=1>.

Woolley, Anita Williams, Christopher F. Chabris, Alex Pentland, Nada Hashmi, and Thomas W. Malone. "Evidence for a Collective Intelligence Factor in the Performance of Human Groups." *Science* 330, no. 6004 (2010): 686–88.

Yue, Fang. "China's OpenClaw Craze: How an Open Source Agent Is Redefining the Future of AI." CEIBS Knowledge, 2026.

Zhu, Gaoxia, Vidya Sudarshan, Jason Fok Kow, and Yew Soon Ong. "Human-Generative AI Collaborative Problem Solving: Who Leads and How Students Perceive the Interactions." arXiv preprint, May 19, 2024. <https://arxiv.org/abs/2405.13048>.

Zuboff, Shoshana. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: PublicAffairs, 2019.

Appendices

Appendix A. Context-maxxing infrastructure building blocks

Table A1. Open-source harnesses

| Harness | Language / runtime | Design emphasis | Typical deployment | MCP support | Notable features / trade-offs |
|--------------------|-----------------------|-----------------------------------------------|-------------------------|-------------|---------------------------------------------------------------------------------------------------|
| OpenClaw | Python | Feature-rich personal agent, big ecosystem | Local or self-hosted | Yes | Very powerful and extensible; heavier RAM use and a larger codebase. |
| NanoClaw | Python + containers | Security-first, containerized sandboxing | Local or server | Yes | Strong isolation; simpler than OpenClaw but with more ops overhead than ultra-light variants. |
| PicoClaw | Python (stripped) | Ultra-light, edge / low-resource hardware | Edge boards, Pi-class | Partial | Very small footprint and fast startup; fewer built-ins, more manual wiring. |
| ZeroClaw | Rust | Performance and portability | Cross-platform binaries | Planned | High throughput and low footprint; smaller ecosystem, more low-level configuration. |
| NullClaw | Zig | Extreme minimalism and speed | Static single binary | No | ~1 MB RAM, sub-2 ms startup, supports many providers; minimal but highly tuned core. |
| Nanobot | Python | Readability, tiny core, research-friendly | Local / Docker | Yes | ~4k LOC core, supports multiple LLM providers and chat channels; good for inspection and hacking. |
| TinyClaw | Mixed (orchestration) | Multi-agent coordination across channels | Bots on chat platforms | Partial | Focuses on orchestrating multiple agents rather than a single "personal AI." |
| OpenHarness | Python | "Reference" harness for learning and research | Local / server | Yes | Demonstrates production patterns (tools, memory, routing) in a clean, didactic codebase. |

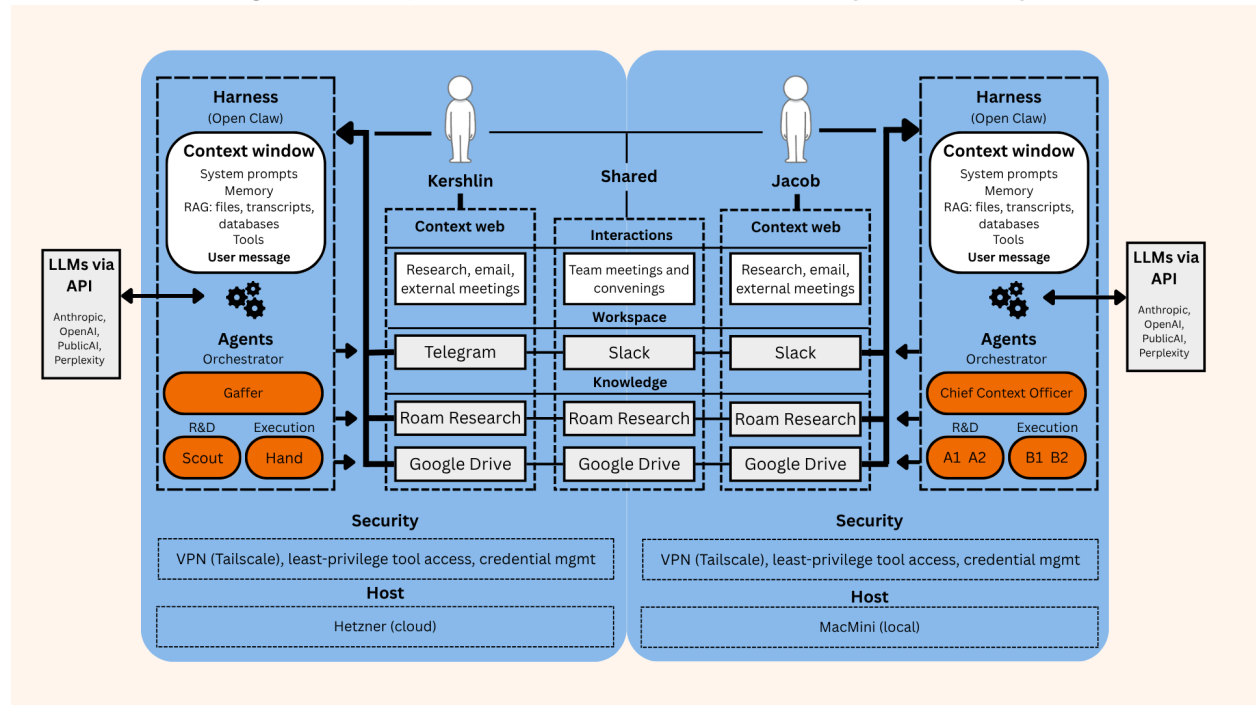
Table 2. LLM options

| LLM category | Typical cost band (input / output) | Example providers and models | How to read this in a harness context |
|----------------------------|------------------------------------------------------------------|----------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| Frontier closed | Roughly \$2–\$6 in / \$12–\$30 out | Anthropic Claude Opus 4.6; OpenAI GPT-5.3/5.4; Google Gemini 3.1 Pro | Use sparingly for hardest reasoning, complex tool orchestration, and high-stakes decisions. |
| Mid-tier closed | Roughly \$0.3–\$1 in / \$1–\$4 out | Claude Sonnet-class; OpenAI “mini” models; Gemini Flash-class | Good default for many agent workflows where quality matters but isn’t safety-critical. |
| Cost-optimized closed | Roughly \$0.02–\$0.1 in / \$0.1–\$0.5 out | OpenAI “nano”/entry-tier chat; smaller Gemini Flash-Lite | Suitable for bulk tasks: tagging, light summarization, simple transformations. |
| Strong open-weight APIs | Roughly \$0.1–\$1 in / \$0.1–\$1 out | DeepSeek V-series; Qwen2.5 Max/72B; Tongyi/Qwen, Hunyuan APIs | Often competitive with mid-tier closed models; good for cost-sensitive but capable workloads. |
| Public / civic / EU-hosted | “Low but variable”; sometimes free tier, sometimes mid-tier-like | PublicAI / Apertus-style services, other civic endpoints | Useful where governance/jurisdiction matters; design around rate limits and variability. |
| Fully open, self-hosted | No API fee; pays in hardware + ops | AI2 OLMo; Apertus weights; other open models on your own GPUs/CPUs | Best when you have steady volume and strong infra; cost per token can be very low at scale. |

Note: Model choice should be guided by the demands of the task, including performance, cost, security, and operational constraints. It is useful to think in tiers defined by both capability and control. At the top are the most advanced proprietary models (for example Claude Opus 4.7, GPT-5.3/5.4, and Gemini 3.1 Pro), which offer the strongest reasoning, coding, and multimodal performance, but at higher cost and with the requirement that sensitive information pass through external infrastructure. A second tier includes lower-cost proprietary models that are well suited to high-volume summarization and straightforward classification (for example, GPT-5.4-mini and -nano, or Claude Sonnet). A third consists of increasingly capable open models delivered through cloud services (including leading Chinese models such as DeepSeek-V3.2 and Qwen2.5 Max), which are becoming more competitive on both price and performance. Beneath these are community-oriented and public-interest models, which place greater weight on transparency, multilingual access, and research or civic purposes than on benchmark leadership (such as PublicAI/Apertus or fully open projects like AI2’s OLMo 2). Proprietary hosted models tend to provide the highest performance, the richest tool ecosystems, and the most reliable service commitments, but they offer less control over data handling. Open models accessed through cloud platforms expand flexibility and can lower costs, though they place more responsibility on the user to manage selection, routing, and quality control. Self-hosting extends that logic further, increasing control and potentially reducing ongoing costs, while also requiring greater technical and operational capacity.

Source: Price Per Token, “LLM API Pricing 2026: Compare 300+ AI Model Costs,” accessed April 11, 2026, <https://pricepertoken.com/>; TLDL, “OpenAI API Pricing 2026 — GPT-5.4, O3, O1 & GPT-4o Cost Per Token,” updated March 5, 2026, accessed April 11, 2026, <https://www.tldl.io/resources/openai-api-pricing>.

Figure A1. Implementation of context-maxxing at Brookings



Note: Kershlin’s agent architecture consists of a single orchestrator (Gaffer) with two specialist sub-agents: one focused on research and development (Scout), the other on analysis and production (Hand). Jacob’s architecture consists of a single orchestrator (Chief Context Officer or CCO) with two sub-agent teams for research and development (A1: Roamer; A2: Locksmith), and drafting and production of outputs (B1: Drafter and B2: Producer).

Appendix B: Context-maxxing competencies

The author's implementation rested on two lightweight templates for specifying relevant context for human-AI workflows. The two templates were designed to work in tandem to produce rich context assets for each Room:

1. **A Policy Innovation Template (PIT).** A set of five questions that helps specify a Room's proposed innovation (see Table B1).
2. **A Policy Critique Template (PCT).** The critique framework applies four evaluative dimensions to a proposed policy innovation (see Table B2).

The two templates function as a matched pair for each Room. Deployed iteratively, these templates helped structure reusable context assets for human-AI workflows. The PIT generated responses using assembled context, while the PCT evaluated and refined those responses.

These assets behaved like a fractal across human-AI workflows (or what the authors have come to term "master keys"): with the help of generative AI's ability to quickly summarize or elaborate, each context asset can be represented as a 250-character tweet, a 250-word summary, a 2.5-page convening brief, or a 25,000-word research report. In practice, Jacob and Kershlin each individually developed longer, "deep research"-length versions of each proposed policy innovation that were stored in their full length in Roam Research and/or Google Drive. 250-word abstracts and 2.5-page summaries of these fuller-length assets were created as more human-readable artifacts for sharing and discussing with each other via Slack.

This allows a single investment in specifying and developing context to return value across multiple audiences (humans and LLMs) and for multiple form factors (assets can be transformed into a range of outputs: draft convening briefs, presentation tools, or drafts of public-facing policy memos).

Table B1. The five-question Policy Innovation Template (PIT)

| Policy innovation template | |
|-------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| Component | Description |
| 1. Specific Problem | What is the problem, articulated in its most specific, irreducible form (root cause rather than symptoms)? |
| 2. Proposed Innovation | What solution to the problem would be transformational rather than incremental, while remaining practical and feasible? |
| 3. Implementation strategy | What concrete mechanism makes the transformation possible? |
| 4. Specific actions | Who needs to do what, by when, and with what accountability? |
| 5. Implications and open questions | What downstream effects, epistemic uncertainties, and learning agenda follow? |

Table B2. The four-question Policy Critique Template (PCT)

| Policy critique template | |
|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| Component | Description |
| 1. Policy nuance | What is the problem, articulated in its most specific, irreducible form (root cause rather than symptoms, at system level)? |
| 2. Political feasibility | Will it resonate with the relevant decisionmakers? Is it realistic about power and interests? |
| 3. Implementation strategy | Can the named actors actually do this, given resources and capacity? |
| 4. Communication strategy | Is it persuasive? Does it cut through? How will it be received by diverse audiences? |

Table B3. Security measures

| Safeguard | What it means | Why it matters for agent harnesses |
|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Private network access | Run the agent on an encrypted private network (mesh VPN) linking laptop and phone, rather than exposing it directly to the public internet. | Reduces the chance that outsiders can find, access, or attack the system remotely. |
| Isolation from other systems | Use containers, virtual machines, firewalls, or local-only access to separate the agent from other devices and services. | Limits the damage if the agent, or one connected tool, is compromised. |
| Secure storage for passwords and keys | Store API keys, passwords, and access tokens in an encrypted vault or other protected location. | Prevents sensitive credentials from being exposed through logs, files, browser sessions, or agent interactions. |
| Clear operating rules | Give the agent standing instructions (system prompts) about what it may and may not do, especially with sensitive information or high-risk actions. | Reduces the risk that malicious or misleading instructions (prompt injection) cause the agent to reveal information, ignore boundaries, or take unsafe actions. |
| Human approval for high-risk actions | Require a person (human-in-the-loop) to approve actions such as sending external emails, deleting files, moving money, or changing permissions. | Prevents the agent from taking consequential actions (sending an email or transferring money) on its own because of an error, malicious instruction, or misunderstanding. |
| Separate accounts for the agent | Give the agent its own accounts, inboxes, folders, or workspaces where appropriate. | Limits blast radius if the agent is compromised and makes the agent easier to monitor, restrict, and shut down without affecting a user's main accounts. |
| Least-privilege access | Minimal, scoped permissions to each system and tool. | Reduces the harm that can result from mistakes, misuse, or security failures. |
| Monitoring and shutdown (kill switch) | Keep logs, watch for unusual activity, and maintain a simple way to pause or disable the agent. | Enables faster detection and response if the system behaves unexpectedly. |

BROOKINGS

1775 Massachusetts Ave NW
Washington, DC 20036
(202) 797-6000
www.brookings.edu