

What Aspects of Tacit Knowledge are Structurally Excluded From Generative AI?: A Conceptual Framework of Mediation, Structure, and Representation

Takashi Onoda¹ and Yasunobu Ito²

¹Organization for Co-Creation Research and Social Contributions, Nagoya Institute of Technology, Gokiso-cho, Showa-ku, Nagoya, Aichi, Japan

²Japan Advanced Institute of Science and Technology, Asahidai 1-1, Nomi-City, Ishikawa, Japan

ABSTRACT

Recent advances in generative AI, particularly large language models and multimodal foundation models, have renewed interest in whether tacit knowledge can be learned or reproduced by machines. While prior studies emphasize the growing ability of generative AI to approximate patterns of human reasoning, judgment, and action, comparatively little attention has been paid to what aspects of tacit knowledge are excluded by design. This paper addresses this gap by asking a conceptual question: which aspects of tacit knowledge are structurally excluded from contemporary generative AI research? Rather than treating tacit knowledge as a single implicit capability, this study reorganizes prior research into three analytical perspectives: mediation, structure, and representation. From this viewpoint, tacit knowledge is sustained by processes that translate practice into communicable forms, by social and cultural structures that stabilize judgment and action, and by representational practices that constitute tacit knowledge as an object of analysis. These perspectives are then used to examine recent developments in generative AI. The analysis shows that current generative AI systems primarily engage with the externalized outcomes of tacit knowledge, such as observable reasoning patterns or action trajectories, while leaving its formative conditions unaddressed. Processes of mediation, social and institutional structures, and reflexive representational practices remain outside model design. These limitations are not merely technical but reflect structural design choices embedded in contemporary AI research. By clarifying these boundaries, this paper provides a conceptual framework for reconsidering the division of roles between human practice and generative AI in future socio-technical systems.

Keywords: Tacit knowledge, Generative AI, Human-centered system design

INTRODUCTION

Recent advances in large language models (LLMs) have transformed modes of knowledge production and rapidly expanded the scope of generative AI. Representational learning frameworks initially developed for language have been extended to images, speech, video, and sensor data, and more recently to multimodal foundation models that integrate perception, language, and action.

Received February 7, 2026; Revised April 2, 2026; Accepted April 21, 2026; Available online July 20, 2026

© 2026 The Authors. This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.

For more information, see <https://creativecommons.org/licenses/by-nc-nd/4.0/>

These developments reflect a growing interest in model architectures that address cognition, judgment, and action in an integrated manner. As a result, the question of how far generative AI can approach human practice has become a concrete research problem. Addressing this question requires attention not only to explicit, describable knowledge, but also to non-explicit forms of understanding commonly referred to as tacit knowledge.

Recent studies report that foundation models can reproduce certain external patterns of tacit knowledge, such as pragmatic language use or skill-like behavior (Open X-Embodiment Collaboration et al., 2023; Ouyang et al., 2022). However, because these models rely primarily on observable data, the conditions under which such practices are constituted remain insufficiently examined. This limitation is evident in discussions of “AI-ready” data, which frame tacit knowledge as something that must be translated into computable forms (Bommasani et al., 2021; Ng, 2022).

Against this background, this paper asks: Which aspects of tacit knowledge have been presupposed and excluded in generative AI research? Drawing on prior work in MOT, knowledge science, and STS, the paper organizes tacit knowledge into three analytical perspectives—mediation, structure, and representation—and uses this framework to examine the scope and limits of generative AI.

LITERATURE REVIEW

Tacit knowledge is commonly understood as knowledge that is difficult to articulate verbally and is grounded in experience and practice. While often described as “implicit” or “subjective,” tacit knowledge has been treated in prior research not merely as unspoken information, but as a foundational condition that enables knowledge production itself.

The conceptual origin of tacit knowledge can be traced to Michael Polanyi’s critique of objectivist epistemology. In *Personal Knowledge*, Polanyi challenged the assumption that scientific knowledge exists as an objective system independent of human actors. Through the proposition that “we know more than we can tell,” he argued that knowing is constituted through personal participation, integrating bodily skills, judgment, and commitment (Polanyi, 1962). From this perspective, tacit knowledge is not simply knowledge that has not yet been verbalized but encompasses bodily and normative conditions that make knowing possible in the first place.

Building on Polanyi’s intervention, subsequent research on tacit knowledge developed along two major trajectories. Science and Technology Studies (STS) focused on the social and cultural foundations of tacit knowledge as embedded in practice, while Management of Technology (MOT) examined how tacit knowledge is transformed, shared, and circulated within organizations. Drawing on these traditions, this paper organizes prior discussions of tacit knowledge into three analytical layers: mediation, structure, and representation.

In the MOT literature, tacit knowledge has been treated as a central issue in organizational learning and capability development, particularly in the context of education, training, and on-the-job learning. Early studies

emphasized that expert intuition and experience function as critical resources for creative problem-solving (Leonard & Sensiper, 1998), while later work examined how experiential accumulation contributes to the development of organizational capabilities (Zollo & Winter, 2002; Teece et al., 1997).

The most influential framework in this domain is the knowledge creation theory proposed by Nonaka and Takeuchi. Their SECI model conceptualizes knowledge creation as a cyclical process consisting of four phases: Socialization, Externalization, Combination, and Internalization (Nonaka & Takeuchi, 2020). Socialization refers to the sharing of bodily, relational, and situational experience through co-presence and observation. Externalization connects such experience to concepts through metaphor, dialogue, and narrative articulation. Combination integrates these concepts into formal knowledge systems such as documents or manuals, while Internalization re-embeds systematized knowledge into practice as new tacit knowledge.

What this model makes explicit is that tacit knowledge does not circulate naturally within organizations. Instead, a series of mediating operations—editing experience, translating it into concepts, reorganizing it as formal knowledge, and reintegrating it into practice—are required. Accordingly, the central concern of MOT-oriented tacit knowledge research lies in the institutional design of these mediating processes.

In contrast to the process-oriented focus of MOT, STS research has examined tacit knowledge by analyzing the social and material arrangements through which scientific and professional practices are organized. Laboratory studies demonstrated that experimental procedures, records, and visualizations rely on tacit understandings, local conventions, and shared expectations (Latour & Woolgar, 1986). Knorr-Cetina further showed that different scientific fields develop distinct epistemic cultures through tacitly shared modes of practice and collaboration (Knorr-Cetina, 1999).

Within this tradition, tacit knowledge is understood not as an internal attribute of individuals, but as knowledge constituted within networks of practice that intertwine instruments, human actors, and organizational arrangements. Building on these insights, Collins provided a systematic formulation of tacit knowledge as a social and intersubjective phenomenon. He distinguished three forms of tacit knowledge: bodily, relational, and collective (Collins, 2010).

Bodily tacit knowledge refers to forms of skill and mastery inscribed through repeated sensorimotor experience and cannot be acquired through rules or manuals alone. Relational tacit knowledge corresponds to unarticulated understandings shaped by organizational culture, institutional arrangements, and local norms, which may remain unspoken due to social constraints or lack of incentive. Collective tacit knowledge is embedded even more deeply in social practice, encompassing linguistic fluency, cultural conventions, evaluative standards, and shared modes of judgment that can only be acquired through participation in a community.

A key implication of Collins's framework is that tacit knowledge is sustained by social structures and cannot be reduced to hidden individual skills. In particular, collective tacit knowledge represents a domain that is least amenable to formalization or mechanical reproduction, even when partial articulation is possible.

The third analytical layer concerns representation: the methodological practices through which tacit knowledge is rendered observable, describable, and analyzable. Representation here does not aim to reproduce tacit knowledge itself, but to construct forms through which the conditions and premises of judgment can be examined.

Engineering-oriented approaches such as knowledge graphs and digital twins exemplify one response to this challenge by mapping practical knowledge into computationally manipulable structures (Ehrlinger & Wöß, 2016; Hogan et al., 2022; Tao et al., 2019). These approaches enable the explicit representation of purposes, constraints, and relationships embedded in practice, and have been applied to partial visualization of tacit expertise in domains such as caregiving and bodily instruction (Ijuin, 2022; Sato et al., 2024). Recent work has further extended this framework toward human-centered digital knowledge twins that support knowledge circulation and reuse (Uchihira et al., 2025).

At the same time, qualitative approaches—particularly ethnography—have long addressed the representation of tacit knowledge by reconstructing interactional patterns and institutional arrangements that make specific judgments and practices possible (Umemoto, 2004; Kosaka, 2007; Miyamoto & Kanda, 2020). Ethnographic analyses of research infrastructures further demonstrate how epistemic cultures shape scientific practice and collaboration (Onoda & Ito, 2022).

Although engineering and qualitative approaches differ in methodological orientation, both function as representational practices that establish tacit knowledge as an object of inquiry. While engineering approaches emphasize operational manipulability, qualitative approaches preserve contextual and relational complexity. Taken together, they provide complementary means of engaging with the multilayered nature of tacit knowledge.

This review confirms that tacit knowledge is not a unitary phenomenon, but a multilayered object spanning mediation, structure, and representation. Mediation concerns the processes through which tacit knowledge is translated and circulated, structure addresses the social and cultural conditions that sustain tacit knowledge, and representation refers to the methodological practices that render tacit knowledge analyzable. The following section examines how generative AI engages with each of these layers.

DISCUSSIONS

In discussions of generative AI, the “incorporation of tacit knowledge” is frequently invoked. However, far less attention has been paid to the conditions under which tacit knowledge itself is constituted, and to which aspects of tacit knowledge are structurally presupposed—and thus excluded—in the design of generative AI systems. To address this gap, this section examines generative AI research through the three analytical perspectives introduced above: mediation, structure, and representation. This framework makes it possible to clarify not only what generative AI can reproduce, but also which conditions it does not treat as design objects.

As reviewed earlier, research in MOT and knowledge science has consistently emphasized that tacit knowledge becomes connected to explicit knowledge only through mediating operations such as editing, translation, and conceptual reorganization. The SECI model proposed by Nonaka and Takeuchi makes clear that the circulation of tacit knowledge as organizational knowledge requires processes of externalization that actively reorganize experience into communicable forms (Nonaka & Takeuchi, 2020).

From the perspective of mediation, it is appropriate to examine language-centered approaches in generative AI, given that judgment, reasoning, and meaning-making are primarily articulated through language. Recent advances include methods that incorporate human judgment and preference as learning signals. Reinforcement Learning from Human Feedback (RLHF), introduced in InstructGPT, uses human evaluations of model outputs as reward signals to align responses with preferred behaviours (Ouyang et al., 2022). Human Preference Learning (HPL) further systematizes this approach by explicitly addressing the partial, context-dependent, and variable nature of human evaluations (Jiang et al., 2024). In addition, Chain-of-Thought (CoT) prompting elicits step-by-step reasoning traces that approximate the external form of expert reasoning (Wei et al., 2022).

However, these approaches learn the outcomes of judgment and reasoning, or their expressed forms, rather than the mediating processes through which information is selected, contextualized, and translated into meaning. This limitation is not merely technical, but structural: generative AI systems are not designed as agents that organize experience or perform editorial work on situated practice. Consequently, while generative AI can reproduce externalized patterns of judgment and reasoning, it does not engage with the mediating processes that constitute tacit knowledge as explicit knowledge in the first place.

As discussed in the review, tacit knowledge is not reducible to individual bodily skill or experience, but is sustained by shared social and cultural conditions, including norms, role expectations, and interactional orders within communities of practice. Collins's framework highlights that tacit knowledge functions only insofar as it is embedded in such structural conditions (Collins, 2010).

To examine this structural dimension, it is necessary to consider generative AI research that addresses action and embodied practice, rather than language alone. Recent work on multimodal and embodied foundation models integrates language, vision, and action into unified architectures. For example, RT-2 adopts a Vision–Language–Action framework that directly links perceptual input and linguistic knowledge to action selection (Brohan et al., 2023). Open X-Embodiment and RT-X further aim to generalize skills across heterogeneous robots by integrating operation logs collected under differing embodiments and environmental constraints (Open X-Embodiment Collaboration et al., 2023).

These studies can be understood as approaching aspects of tacit knowledge related to bodily skill and the coupling of situational judgment and action. At the same time, the social and cultural conditions that define why particular actions are regarded as appropriate—such as normative expectations, role relations, and shared conventions—are not represented as components of the

model's world. As a result, generative AI can reproduce patterns of behaviour, but cannot assess their social meaning or contextual appropriateness. This absence of structural conditions as design objects constitutes a fundamental limitation that persists regardless of modality or architectural sophistication.

The third analytical perspective concerns representation: the practices through which tacit knowledge is rendered observable and analyzable. As demonstrated by both engineering approaches, such as knowledge graphs and digital twins, and qualitative methodologies, such as ethnography, engaging with tacit knowledge requires an explicit observational standpoint that reconstructs social structure and mediating processes into analyzable forms.

Contemporary generative AI systems lack such representational practices internally. In RLHF and HPL, models learn from evaluative outcomes, but do not observe or describe the premises, interactional contexts, or normative assumptions underlying those evaluations. Similarly, while CoT enables the generation of reasoning explanations, it does not provide an observer's perspective capable of reconstructing the conditions under which those explanations become meaningful.

The same limitation applies to speech and vision–action models. Although these systems reproduce external forms of speech and movement, they do not perform representational operations that extract tacit elements such as etiquette, role relations, or shared purposes as analyzable objects. Consequently, generative AI reconstructs the outward patterns of practice without access to the premises that sustain them.

Taken together, this analysis shows that generative AI has achieved notable success in reproducing external manifestations of tacit knowledge, while leaving its constitutive conditions outside the scope of system design. From the perspective of mediation, editorial and translational processes are not treated as learning targets. From the perspective of structure, social and cultural preconditions are not represented as elements of the model's world. From the perspective of representation, generative AI lacks an observational standpoint that would allow tacit knowledge to be reconstructed as an analyzable object.

Importantly, these limitations should not be interpreted as simple failures of capability. Rather, they reflect the current design assumptions of generative AI. By clarifying these assumptions, the mediation–structure–representation framework provides a basis for more precise discussions about which aspects of human practice can be meaningfully connected to generative AI, and which remain dependent on human interpretation and methodological intervention.

CONCLUSION

This paper addressed the research question of which aspects of tacit knowledge have been presupposed and excluded in generative AI research. Rather than treating tacit knowledge as a single, homogeneous property, the paper reconfigured it through three analytical perspectives—mediation, structure, and representation—drawing on established discussions in

tacit knowledge research. This framework provides a reference point for systematically comparing generative AI research with prior theoretical work.

The central contribution of this paper lies in clarifying that, before debating the reproducibility or performance limits of tacit knowledge in generative AI, it is necessary to distinguish which underlying conditions of tacit knowledge are not incorporated into current AI design assumptions. From this perspective, the limitations of generative AI are not merely matters of capability but are rooted in a design space that presupposes specific exclusions.

This argument does not deny the significant achievements of recent generative AI. On the contrary, the ability of large language models and multimodal foundation models to reproduce externalized patterns of human reasoning, judgment, and action has made the gap between AI design assumptions and long-standing concerns in knowledge science and Science and Technology Studies increasingly visible.

Accordingly, the three-layer framework proposed here serves not as a final verdict on the limits of generative AI, but as a conceptual scaffold for examining at which levels generative AI can be meaningfully connected to human practice, and which levels remain dependent on human participation and methodological intervention. The contribution of this paper thus lies in providing a basis for rethinking how collaboration between humans and generative AI can—and should—be structured going forward.

ACKNOWLEDGMENT

This work was supported by JSPS KAKENHI Grant Number 23K12517.

REFERENCES

- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., et al. (2021). On the opportunities and risks of foundation models. arXiv preprint arXiv:2108.07258.
- Brohan, A., Brown, N., Carbajal, J., Chebotar, Y., Dabis, J., Finn, C., et al. (2023). RT-2: Vision-language-action models transfer web knowledge to robotic control. arXiv preprint arXiv:2307.15818.
- Collins, H. (2010). Tacit and explicit knowledge. University of Chicago Press, Chicago.
- Ehrlinger, L., and Wöß, W. (2016). Towards a definition of knowledge graphs. In Proceedings of SEMANTiCS 2016, Leipzig, Germany.
- Hogan, A., Blomqvist, E., Cochez, M., d'Amato, C., de Melo, G., Gutierrez, C., et al. (2022). Knowledge graphs. ACM Computing Surveys, 54(4), 1–37.
- Knorr-Cetina, K. D. (1999). Epistemic cultures: How the sciences make knowledge. Harvard University Press, Cambridge, MA.
- Latour, B., and Woolgar, S. (1986). Laboratory life: The construction of scientific facts. Princeton University Press, Princeton, NJ.
- Leonard, D., and Sensiper, S. (1998). The role of tacit knowledge in group innovation. California Management Review, 40(3), 112–132.
- Nonaka, I., and Takeuchi, H. (2020). The knowledge-creating company (new ed.). Toyo Keizai Inc., Tokyo. (Originally published in English, 1995).

- Onoda, T., and Ito, Y. (2022). On the boundary of services and research collaborations in Japanese state-of-the-art academic research infrastructures. *Science and Public Policy*, 49(3), 488–498.
- Open X-Embodiment Collaboration, Brohan, A., Chebotar, Y., Finn, C., Geng, D., Hausman, K., et al. (2023). Open X-embodiment: Robotic learning datasets and RT-X models. arXiv preprint arXiv:2310.08864.
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., et al. (2022). Training language models to follow instructions with human feedback. arXiv preprint arXiv:2203.02155.
- Polanyi, M. (1962). *Personal knowledge: Towards a post-critical philosophy*. University of Chicago Press, Chicago.
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., et al. (2022). Chain-of-thought prompting elicits reasoning in large language models. arXiv preprint arXiv:2201.11903.
- Zollo, M., and Winter, S. G. (2002). Deliberate learning and the evolution of dynamic capabilities. *Organization Science*, 13(3), 339–351.