Generative behaviors as key targets for cognitive models

Judith E. Fan
Department of Psychology, Stanford University

Author Note

Correspondence should be addressed to Judith E. Fan, 450 Jane Stanford Way, Building 420, Stanford, CA, USA 94305. Email: jefan@stanford.edu. This work was supported by NSF CAREER #2436199, NSF DRL #2400471, a Hoffman-Yee award from the Stanford Human-Centered AI Institute, and seed grants from the Stanford Accelerator for Learning. Thanks to Daniel Yamins, Robert Hawkins, Mike Frank, and Lio Wong for helpful discussions.

Abstract

Human behavior is fundamentally generative: people create pictures, write stories, compose music, and engage in conversation. Traditional approaches in psychology and cognitive science have not focused on this open-endedness, instead favoring more constrained task settings that admit a limited set of outcomes. While those approaches have been fruitful, new approaches might be needed to develop a unified understanding of the generative, open-ended behaviors that are so emblematic of human cognition. This paper argues for the value of generative behaviors as targets for cognitive modeling by providing rich behavioral data that reveal how multiple cognitive processes coordinate. Drawing production serves as a case study illustrating this approach, showing how perception, memory, social inference, and motor control coordinate flexibly based on communicative context. Recent advances in generative AI offer both new tools for modeling open-ended human behavior and new comparative targets for understanding similarities and differences between human and machine intelligence. However, applying these tools effectively might require new experimental paradigms, larger datasets, and careful consideration of what mechanistic correspondence between models and human cognition is necessary for scientific progress. Embracing the open-ended nature of human thought and behavior poses methodological challenges but offers a promising path toward understanding the most distinctive aspects of human intelligence.

Keywords: artificial intelligence, behavioral measurement, computational modeling, drawing, production

The technological advances powering recent progress in AI have accelerated the exploration of many pathways to intelligence. These advances stand to challenge traditional accounts of human exceptionalism when it comes to a variety of cognitive capacities, including the ability to perform complex reasoning tasks and use language in natural-sounding ways (Bommasani et al., 2021; Demszky et al., 2023; Frank and Goodman, 2025; Binz et al., 2025). In light of these developments, what does it look like to make continued progress in our understanding of human cognition within the broader landscape of possible "solutions" to the full suite of cognitive challenges that have been so far posed by nature?

In this paper, I will argue that further progress can be made by prioritizing the study of generative behaviors—that is, behavior approximating the complexity and open-endedness that humans display in real-world settings (Carvalho & Lampinen, 2025). The first section will consider what generative behaviors are and why the resulting data offer valuable targets for cognitive modeling efforts that complement what their discriminative counterparts typically provide. The second section will review a case study that exemplifies this approach, namely the use of drawing production tasks to gain understanding of how multiple cognitive processes interact to support learning and communication. The third section will explore the relationship between the study of generative behaviors produced by humans and those produced by generative AI systems. I will then close by noting further opportunities and challenges presented by the study of generative behaviors.

What are generative behaviors and why study them?

The scientific study of human cognition is guided by the goal of making progress along two dimensions: The first is to reveal the mind's functional architecture—what the parts are, what they are for, and how these parts are connected to each other. The second is to explain behavior—to accurately predict what people think, feel, do, and say in a variety of situations. Good behavioral tasks are necessary for making scientific progress along both dimensions because behavior is what can be observed and measured[1], while everything else that scientists might seek to learn about how the mind is organized must be inferred from those measurements. But what kinds of behavioral tasks are most valuable for making progress?

Behavioral tasks can be characterized along a continuum according to the nature of the responses they require. At one end are “discriminative” tasks, which involve selection or judgment over a limited set of options. These include making ratings along a fixed scale, classifying stimuli into predetermined categories, and producing estimates along specific feature dimensions. At the other end are “generative” tasks, which involve producing complex, structured outputs from an effectively infinite space of possibilities. These include drawing pictures, composing explanations, and solving open-ended problems. For instance, when studying a person’s understanding of a visual concept (e.g., *dog, chair, truck*), this individual could be asked to select which of a few different images matches that concept (discriminative), or they could draw an image that evokes that concept (generative). Most behavioral tasks that are used in cognitive psychology fall somewhere along this continuum (Figure 1).

[1] In principle, we can adopt a fairly expansive notion of what "behavior" is and include measurement of the behavior of neurons and the behavior of the circuits they comprise in the brain. Sometimes those measurements can be useful for distinguishing among possible internal mechanisms. But by and large, the behavior of the whole organism—in particular, an individual human—is the focus of this paper, because it is what we take the behavior of the neurons to be there to support.

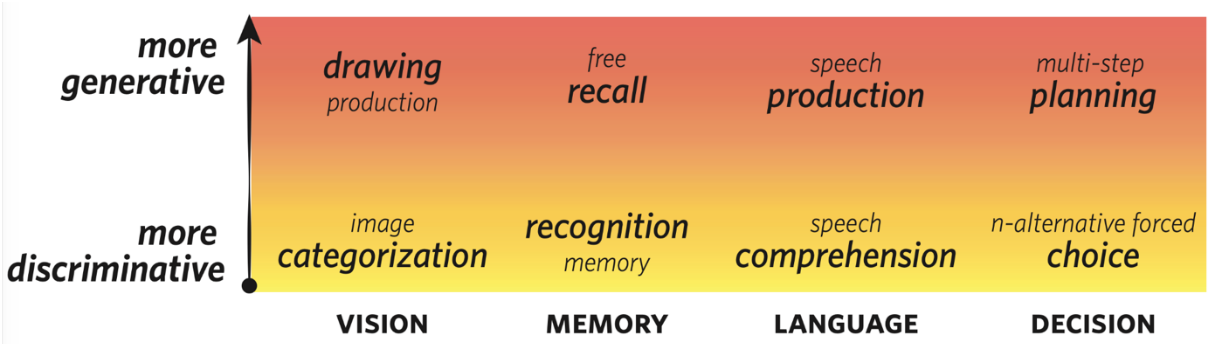

Figure 1: Examples of generative tasks and discriminative tasks that have been used to investigate various aspects of human cognition (e.g., vision, memory, language, and decision-making). There is no sharp boundary distinguishing generative from discriminative tasks. Instead, most of these behavioral task paradigms are taken to vary along a continuum between more generative variants and more discriminative variants.

There are several hallmarks of behavioral tasks that have long been considered to be valuable in psychology: First, good behavioral tasks enable accurate and precise measurement. For example, there are some tasks that people can perform repeatedly, giving researchers the ability to use repeated-measures designs to obtain more precise estimates at the individual level. Second, good behavioral tasks allow for measurement of sufficient behavioral variability that there is something to explain, but not so much variation that it would be impractical to collect enough data to explain it. Third, good behavioral tasks enable experimental manipulations, so that researchers can infer causal relationships between the variables that were manipulated and the variables that were measured. Discriminative tasks have been particularly successful in satisfying these criteria, enabling researchers to isolate and characterize elementary cognitive processes, such as selective attention and working memory.

Generative behavioral tasks offer an additional virtue that has only recently come to be appreciated: they provide finer-grained information about the mental states that give rise to behavior. When two people produce the same response when performing a discriminative task—

for example, both selecting "Somewhat agree" when answering a survey question—it can be difficult to determine whether they arrived at that response via the same mental processes, because many different mental states can produce the same observed response. By contrast, generative outputs leave behind behavioral traces that vary along many dimensions simultaneously, providing more potential avenues for distinguishing between different underlying mental states, provided one has the appropriate tools to measure and model that variation. When two people draw a dog, the resulting drawing reveals not only whether they possess the concept of a dog, but also which features come to mind. The challenge of measuring and modeling behavioral variation along many dimensions simultaneously has historically limited the use of generative tasks. But the widespread adoption of online data collection and machine learning methods in psychology have made such work much more viable in recent years (Fan et al., 2023).

Generative tasks are also valuable for advancing understanding of how multiple cognitive processes interact. For instance, producing a coherent explanation of a scientific concept—e.g., why ice floats in water—requires retrieving the most relevant pieces of knowledge about water molecules and density, understanding the causal relationship between water's molecular structure and its physical properties, selecting appropriate language to express that relationship, and judging whether the explanation achieves its communicative goal. Cognitive models capable of explaining such behavior—computational instantiations of psychological theories that specify causal relationships between mental representations—must therefore specify not only what these component processes are, but also how they interact with each other (Newell, 1973). This requirement makes generative behaviors particularly useful for testing hypotheses about how cognitive processes work in concert.

The study of generative behaviors complements the insights gained by using discriminative tasks to target individual cognitive processes. A common strategy is to develop models of each process separately (e.g., perception, causal inference, decision-making), with the expectation that these models can eventually be integrated to explain more complex behavioral phenomena. Whether this strategy succeeds depends on the nature of the interactions between processes. When cognitive processes interact in straightforward ways—for example, operating independently or through simple additive contributions—models of those individual processes may compose gracefully to predict more complex behavior. However, when processes interact in more complex ways that go beyond independent or linear contributions, models of isolated processes may fail to predict behavior when multiple processes must be coordinated—even if those models accurately capture individual processes (Anzellotti & Coutanche, 2018; Almaatouq et al., 2024). Because generative behaviors inherently require the coordination of multiple processes, they provide crucial tests of whether cognitive models of individual processes compose successfully (Figure 2). However, to test these models rigorously might also require substantially larger datasets and compute budgets than have historically been used in psychology and cognitive science.

Both discriminative and generative tasks have important roles to play, but generative behaviors remain underutilized sources of evidence for testing cognitive models, particularly when the goal is to understand how multiple processes interact to produce complex behavior. The next section focuses on drawing production as a case study for this approach. Moreover, modern AI systems are themselves now capable of producing drawings, language, and other complex outputs, making it particularly timely to develop rigorous methods for studying and modeling generative behaviors in both humans and machines.

Insights from drawing production as a case study

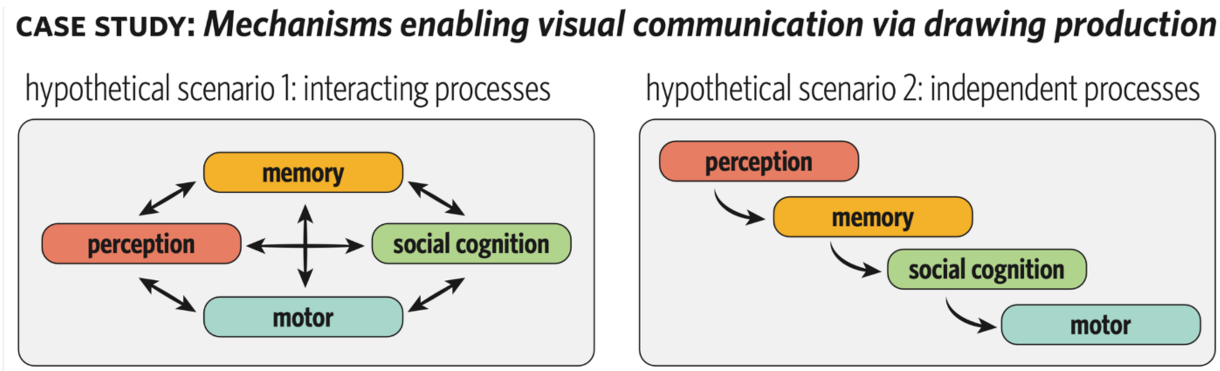

Figure 2: Schematic illustrating two hypothetical scenarios for how the cognitive processes that enable visual communication through drawing production might be organized. In one scenario, there are substantial interactions among the relevant cognitive processes (i.e., perception, memory, social cognition, and motor planning). In the other scenario, these same processes operate independently and serially. This figure is not comprehensive: there are other possible scenarios that are intermediate to these two with respect to the degree of interaction. Studying generative behaviors, such as drawing production, can provide crucial data to help adjudicate among these scenarios.

Drawing production—the act of creating images to communicate ideas—represents a particularly rich domain for investigating how multiple cognitive processes interact to give rise to generative behavior. Drawings can vary tremendously in their visual properties, from highly realistic illustrations to schematic diagrams, and prior work suggests that this variation reflects differences in the balance of contributions from perception, memory, and social inference recruited during production (Fan et al., 2023). By examining how people produce drawings under different conditions, researchers can gain insight into both the content of underlying mental representations and the dynamics of how different cognitive processes coordinate with each other. Three insights from research on drawing production illustrate these contributions.

First, drawing tasks can provide a way to probe mental representations that complements what can be learned from discriminative tasks. When someone draws a dog, the resulting image can reveal what features they consider most important (e.g., four legs, tail, ears, snout), how those features are arranged in space, and what level of detail they choose to preserve (Fan et al., 2018; Yang & Fan, 2021; Lu et al., 2023). Discriminative tasks require researchers to make theory-laden choices about which dimensions to probe, potentially missing features that are important but unanticipated. However, drawing tasks have their own limitations: a participant might fail to include an important feature because they cannot recall it or cannot successfully convey it through drawing. The value lies in the complementary nature of these approaches—visual discrimination tasks provide more direct measurement of specific dimensions but are vulnerable to experimenter bias in selecting response options, while drawing production tasks can reveal unanticipated aspects of mental representations but are constrained by what participants can retrieve and act upon.

Second, studying drawing production can help to reveal the iterative and interactive nature of cognitive processes during complex behaviors that unfold over time. Creating an observational drawing might seem like it could proceed through a straightforward sequence: perceive the scene, form a mental representation of it, then execute a motor plan to reproduce its appearance. However, prior work using eye-tracking suggests that it might not be so simple. In one study, participants continually moved their eyes between the source and their drawing throughout drawing production (Tchalenko & Miall, 2009). Moreover, when participants drew without being able to refer to the source or to their own drawing-in-progress, the resulting drawings exhibited substantial distortions. These findings suggest that drawing relies on ongoing coordination between perception (of both the source and drawing-in-progress), memory (for maintaining task-relevant information), and action (for producing marks), rather than on a simple linear progression from perception to action. However, the exact nature of that coordination requires further empirical investigation. Recent computational work has begun to model the sequential nature of sketch generation (Vinker et al., 2025), which might help to formalize different hypotheses about how these processes interact over time.

A third insight emerging from recent work is that the relative contributions of different cognitive processes during drawing production appear to vary depending on context and goals, resulting in systematic differences in how drawings look—from highly realistic depictions to sparse, schematic abstractions. This variation relates to a longstanding theoretical puzzle about how drawings derive their meaning: do they resemble the entities they depict (depending primarily on visual processing for interpretation) or function as symbols (whose meanings are determined by learned associations)? Recent empirical work suggests that drawings may span a continuum between these poles depending on the conditions under which they were produced (Yang & Fan, 2021). When perception dominates—as when creating an illustration of a visible scene—the resulting drawings can be interpreted with minimal reliance on prior experience. When memory dominates—as when drawing something one has previously seen—drawings can exhibit systematic distortions consistent with memory's reconstructive nature, such as boundary expansion or inclusion of details that were not actually present but are semantically related (Bainbridge & Baker, 2020). When social inference plays a key role—as when communicating with a particular viewer—drawings may become more abstracted, with studies showing that repeated communication leads to sparser drawings that nevertheless remain effective for individuals sharing interaction history (Hawkins et al., 2023). When drawings are used to convey even more abstract information, such as numerical concepts (Holt et al., 2024) or causal mechanisms (Huey et al., 2023), the balance shifts even further towards reliance on prior knowledge and away from perceptual fidelity.

Several important questions would be valuable to investigate in future work. What computational mechanisms account for the flexible coordination between cognitive processes (e.g., perception, memory, planning, action selection) to support drawing production in this wide variety of settings? What principles apply only to drawing production, and which ones also govern generative behaviors in other domains, including language production (Levelt, 1989; Gleitman et al., 2007), gesture (Goldin-Meadow, 1999), explanation (Chi et al., 1994), physical assembly (McCarthy et al., 2023; Binder et al., 2025), gameplay (van Opheusden et al., 2023; Allen et al., 2024; Chu et al., 2025), and constructive activities in educational contexts (Papert, 1980; Chi & Wylie, 2014)? It is plausible that all of these generative behaviors might need to overcome some of the same abstract computational challenges, such as the encoding of sensory inputs into a task representation that enables the efficient selection of goal-relevant actions. But the content of these task representations and which actions are available might differ substantially across domains.

Generative behaviors in humans and machines

Recent advances in artificial intelligence have produced systems that are themselves capable of generative behaviors. Large language models generate coherent text, diffusion models produce images, and multimodal systems can create diagrams that blend visual and text elements. These machine-learning based AI systems exhibit a level of open-ended expressivity that is reminiscent of human generative behaviors, producing outputs that are not limited to selecting from a narrow set of predetermined options. These new capabilities present an opportunity to study generative behaviors in both humans and machines, and to understand the similarities and differences in how they produce such outputs.

The fact that AI systems can now engage in generative behaviors provides an important existence proof: computational artifacts are capable of approaching human-level behavioral expressivity in open-ended tasks. Having such artifacts makes the scientific study of human generative behaviors more tractable, by providing both new analytical tools for characterizing human behavior and new comparative targets against which human behavioral patterns can be evaluated. For example, generative AI systems that can produce sketches could be used to model the computational processes involved in human drawing production, because the images produced by those same systems provide a concrete basis for comparing how humans and AI systems approach the same generative task (Vinker et al., 2022; Mukherjee et al., 2023).

When AI systems are adapted to model human generative behaviors, important questions arise concerning what we learn when those models succeed or fail. If a model achieves high predictive accuracy, when does this imply that similar mechanisms are also engaged by the human mind? In some cases, the functional demands of these generative tasks might constrain the space of viable computational solutions so strongly that different modeling approaches must converge on similar mechanisms to succeed (Cao & Yamins, 2024b). In other cases, multiple systems employing very different internal mechanisms might produce comparable patterns of behavioral outputs. Identifying cases where this second possibility applies is theoretically meaningful: it implies that the constraints used in developing that model (e.g., training data, model architecture, learning objective) are not sufficient to uniquely determine what mechanisms enable some behavior, and that human cognition reflects additional evolutionary and developmental constraints. Developing a wide variety of computational models is valuable for mapping the space of solutions capable of human-like generative behavior and situating human cognition within that space.

These considerations raise a fundamental question about the goals of cognitive modeling: what level of mechanistic correspondence between model and human is necessary for that model to be scientifically useful? The purpose of building cognitive models is not to achieve perfect predictive accuracy—a goal that would ultimately require replicating every detail of human neural architecture. Rather, the purpose is to provide vehicles for formulating and testing alternative theoretical accounts of how cognitive processes are causally related to each other and give rise to behavior. Such models should support both prediction and control: the ability to forecast behavior under novel conditions and to suggest interventions that would change behavior in systematic ways, with a degree of precision that is scientifically and practically meaningful, but no more.

In summary, the emergence of AI systems capable of generative behaviors create new opportunities for psychology and cognitive science. In principle, researchers can adapt these AI systems to model human generative behaviors, starting with systematic comparison between human and AI responses on the same generative tasks (Frank, 2023). However, to fulfill their potential as cognitive models, it will also be necessary to establish clear model-mechanism mappings—determining which components of AI systems correspond to which cognitive processes in humans (Frank & Goodman, 2025). Establishing these mappings remains methodologically challenging and represents an important frontier for this work. Systematic study of human generative behaviors provides a crucial foundation for overcoming these challenges, towards developing unified cognitive models that can explain both how cognitive processes interact in the human mind and what computational principles underlie generative behaviors in both humans and artificial systems.

Summary and future directions

Human cognition is fundamentally generative. People do not merely perceive the world as it already exists but also construct new ways of categorizing the entities around them, generate novel solutions to problems, and invent technologies that extend their cognitive capabilities. These generative capacities have been the engine of cumulative culture: art, music, written language, mathematics are all expressive modalities that humans created and that continue to evolve. This paper has argued that generative behaviors offer a valuable window into these quintessentially human cognitive capacities. Generative tasks, while historically challenging to use in rigorous experimental work, are increasingly feasible to use and provide crucial complements to discriminative tasks for building and testing cognitive models.

Recent advances in machine learning and AI have yielded systems far more capable of generative behaviors than ever before, presenting both challenges and opportunities for cognitive science and psychology (Bommasani et al., 2021; Frank & Goodman, 2025). However, leveraging these systems to build scientific models of the mind will likely require new approaches to experimental design, the collection of larger datasets, major investments in computing resources, and new ways of thinking about mechanistic abstraction in cognitive models (Cao & Yamins, 2024a). Moreover, it is not clear that current generative AI systems are sufficiently capable for these purposes—that they can engage in the kind of continual learning and adaptation needed for them to provide useful models of how humans not only use their existing repertoire of generative behaviors but create new ones. Embracing these more creative aspects of human thought and expression might thus offer a promising path towards advancing theories of distinctly human intelligence.

Recommended Reading

1. Almaatouq, A., Griffiths, T. L., Suchow, J. W., Whiting, M. E., Evans, J., & Watts, D. J. (2024). Beyond playing 20 questions with nature: Integrative experiment design in the social and behavioral sciences. *Behavioral and Brain Sciences*, 47, e33.
2. Fan, J. E., Bainbridge, W. A., Chamberlain, R., & Wammes, J. D. (2023). Drawing as a versatile cognitive tool. *Nature Reviews Psychology*, 2(9), 556-568.
3. Cao, R., & Yamins, D. (2024b). Explanatory models in neuroscience, Part 2: Functional intelligibility and the contravariance principle. *Cognitive Systems Research*, 85, 101200.
4. Carvalho, W., & Lampinen, A. (2025). Naturalistic Computational Cognitive Science: Towards generalizable models and theories that capture the full range of natural behavior. arXiv preprint arXiv:2502.20349.
5. Frank, M. C., & Goodman, N. D. (2025). Cognitive modeling using artificial intelligence. Annual Review of Psychology, 77.

References

Allen, K., Brändle, F., Botvinick, M., Fan, J. E., Gershman, S. J., Gopnik, A., ... & Schulz, E. (2024). Using games to understand the mind. *Nature Human Behaviour*, 8(6), 1035-1043.

Almaatouq, A., Griffiths, T. L., Suchow, J. W., Whiting, M. E., Evans, J., & Watts, D. J. (2024). Beyond playing 20 questions with nature: Integrative experiment design in the social and behavioral sciences. *Behavioral and Brain Sciences*, 47, e33.

Anzellotti, S., & Coutanche, M. N. (2018). Beyond functional connectivity: investigating networks of multivariate representations. *Trends in Cognitive Sciences*, 22(3), 258-269.

Binder, F. J., Mattar, M. G., J. Kirsh, D., & Fan, J. E. (2025). Humans Select Subgoals That Balance Immediate and Future Cognitive Costs During Physical Assembly. Cognitive Science, 49(11), e70135.

Binz, M., Akata, E., Bethge, M., Brändle, F., Callaway, F., Coda-Forno, J., ... & Schulz, E. (2025). A foundation model to predict and capture human cognition. *Nature*, 1-8.

Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., ... & Liang, P. (2021). On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258.*

Cao, R., & Yamins, D. (2024a). Explanatory models in neuroscience, Part 1: Taking mechanistic abstraction seriously. *Cognitive Systems Research*, 87, 101244.

Cao, R., & Yamins, D. (2024b). Explanatory models in neuroscience, Part 2: Functional intelligibility and the contravariance principle. *Cognitive Systems Research*, 85, 101200.

Carvalho, W., & Lampinen, A. (2025). Naturalistic Computational Cognitive Science: Towards generalizable models and theories that capture the full range of natural behavior. *arXiv preprint arXiv:*2502.20349.

Chi, M. T., De Leeuw, N., Chiu, M. H., & LaVancher, C. (1994). Eliciting self-explanations improves understanding. *Cognitive Science*, 18(3), 439-477.

Chi, M. T., & Wylie, R. (2014). The ICAP framework: Linking cognitive engagement to active learning outcomes. *Educational Psychologist*, *49*(4), 219-243.

Chu, J., Zheng, K., & Fan, J. E. (2025). What makes people think a puzzle is fun to solve?. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 47).

Demszky, D., Yang, D., Yeager, D. S., Bryan, C. J., Clapper, M., Chandhok, S., ... & Pennebaker, J. W. (2023). Using large language models in psychology. *Nature Reviews Psychology*, 2(11), 688-701.

Fan, J. E., Yamins, D. L., & Turk-Browne, N. B. (2018). Common object representations for visual production and recognition. *Cognitive Science*, 42(8), 2670-2698.

Fan, J. E., Hawkins, R. D., Wu, M., & Goodman, N. D. (2020). Pragmatic inference and visual abstraction enable contextual flexibility during visual communication. *Computational Brain & Behavior*, 3(1), 86-101.

Fan, J. E., Bainbridge, W. A., Chamberlain, R., & Wammes, J. D. (2023). Drawing as a versatile cognitive tool. *Nature Reviews Psychology*, 2(9), 556-568.

Frank, M. C. (2023). Baby steps in evaluating the capacities of large language models. *Nature Reviews Psychology*, 2(8), 451-452.

Frank, M. C., & Goodman, N. D. (2025). Cognitive modeling using artificial intelligence. *Annual Review of Psychology*, 77.

Gleitman, L. R., January, D., Nappa, R., & Trueswell, J. C. (2007). On the give and take between event apprehension and utterance formulation. *Journal of Memory and Language*, 57(4), 544-569.

Goldin-Meadow, S. (1999). The role of gesture in communication and thinking. *Trends in Cognitive Sciences*, 3(11), 419-429.

Hawkins, R. D., Sano, M., Goodman, N. D., & Fan, J. E. (2023). Visual resemblance and interaction history jointly constrain pictorial meaning. *Nature Communications*, 14(1), 2199.

Holt, S., Fan, J. E., & Barner, D. (2024). Creating ad hoc graphical representations of number. *Cognition*, 242, 105665.

Huey, H., Lu, X., Walker, C. M., & Fan, J. E. (2023). Visual explanations prioritize functional properties at the expense of visual fidelity. *Cognition*, 236, 105414.

Levelt, W. J. (1989). Speaking: From intention to articulation. MIT Press.

Lu, X., Wang, X., & Fan, J. E. (2023). Learning dense correspondences between photos and sketches. *International Conference on Machine Learning* (pp. 22899-22916). *PMLR*.

McCarthy, W. P., Kirsh, D., & Fan, J. E. (2023). Consistency and variation in reasoning about physical assembly. Cognitive Science, 47(12), e13397.

Mukherjee, K., Huey, H., Lu, X., Vinker, Y., Aguina-Kang, R., Shamir, A., & Fan, J. (2023). SEVA: Leveraging sketches to evaluate alignment between human and machine visual abstraction. *Advances in Neural Information Processing Systems*, 36, 67138-67155.

Mukherjee, K., Huey, H., Stoinski, L., Hebart, M., Fan, J. E., and Bainbridge, W. (2025). Drawings of THINGS: A large-scale drawing dataset of 1,854 object concepts. *Behavior Research Methods.*

Newell, A. (1973). You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In *Machine Intelligence* (pp. 121-146). Routledge.

Papert, S. (1980). Mindstorms: children, computers, and powerful ideas. Basic Books, Inc., USA.

van Opheusden, B., Kuperwajs, I., Galbiati, G., Bnaya, Z., Li, Y., & Ma, W. J. (2023). Expertise increases planning depth in human gameplay. *Nature*, 618(7967), 1000-1005.

Vinker, Y., Pajouheshgar, E., Bo, J. Y., Bachmann, R. C., Bermano, A. H., Cohen-Or, D., ... & Shamir, A. (2022). CLIPasso: Semantically-aware object sketching. *ACM Transactions on Graphics (TOG)*, 41(4), 1-11.

Vinker, Y., Shaham, T. R., Zheng, K., Zhao, A., Fan, J. E., & Torralba, A. (2025). Sketchagent: Language-driven sequential sketch generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference* (pp. 23355-23368).

Yang, J., & Fan, J. E. (2021). Visual communication of object concepts at different levels of abstraction. *arXiv preprint arXiv:2106.02775.*