From Images to Symbols: Drawing as a Window into the Mind

Kushin Mukherjee[1], Holly Huey[2], Timothy T. Rogers[1], and Judith E. Fan[2]

[1]University of Wisconsin-Madison, [2]University of California, San Diego

Abstract

Drawing is a powerful cognitive technology for creating external representations of thought. While drawings have long provided inspiration to researchers in many areas of cognitive science, including psychology, artificial intelligence, and neuroscience, these communities have not generally had opportunities to interact and share insights. The goal of this workshop is to bring together perspectives from multiple disciplines to explore the question of how humans use drawings to communicate knowledge, catalyzing new opportunities for multidisciplinary collaboration. We are introducing a novel "flipped" format wherein we will hold three virtual seminars in the weeks leading up to CogSci 2022, each highlighting insights from Machine Learning, Neuroscience, and Developmental Science, respectively. Holding these thematic seminars in advance will enable us to attract a broader audience for our event and focus on promoting informal interaction among workshop attendees at the on-site event.

Keywords: perception, production, concepts, abstraction, learning

Overview and Motivation

One of the most distinctive aspects of human communication is that it goes beyond vocal production — humans have devised many ways to make their ideas both visible and durable over longer timescales relative to the spoken word. From etchings on cave walls to modern digital displays, some of the most significant inventions in human history include technologies that externalize our thoughts in visual form. Perhaps the most basic and versatile of these technologies is drawing. The earliest drawings predate the invention of writing by thousands of years and drawn images have been produced by people from many different cultures (Gombrich, 1989; Clottes, 2008).

A growing body of work has used drawing behavior to shed light on a wide range of cognitive phenomena, including perception/attention (Chamberlain & Wagemans, 2015; Perdreau & Cavanagh, 2014), learning (Fiorella & Zhang, 2018; Fan, Yamins, & Turk-Browne, 2018; Chamberlain, 2018), development (Dillon, 2021; Long, Fan, Chai, & Frank, 2021; Kellogg, 1969; Karmiloff-Smith, 1990; Phillips, Hobbs, & Pratt, 1978), memory (Wammes, Meade, & Fernandes, 2016; Bainbridge, Hall, & Baker, 2019), concepts (Bozeat et al., 2003; Mukherjee, Hawkins, & Fan, 2019; Yang & Fan, 2021), and communication (Fan, Hawkins, Wu, & Goodman, 2020; Hawkins, Sano, Goodman, & Fan, 2021; Huey, Walker, & Fan, 2021).

Despite the fact that humans can effortlessly produce and interpret drawings, their status as being both image-like and symbol-like has posed persistent challenges for existing accounts of how they encode information (Goodman, 1976; Abell, 2009; Greenberg, 2021; Hertzmann, 2020). How can we make progress towards more unified theories that explain how the human mind is capable of making sense of such a broad array of visual inputs, including photos, sketches, diagrams, maps, graphs, text, paintings, and cartoons as well as generate them?

Goals and Approach

The goal of this workshop is to bring together perspectives on this question from multiple disciplines to share insights and formulate strategies for confronting outstanding challenges. Specifically, we will focus on perspectives from Artificial Intelligence, Neuroscience, and Developmental Science, owing to recent empirical advances in each of these fields. First, increasingly mature machine learning frameworks and crowdsourcing technologies have led to the development of algorithms that display impressive sketch understanding and generation capabilities (Ha & Eck, 2017; Jongejan, Rowley, Kawashima, Kim, & Fox-Gieg, 2017; Bhunia et al., 2020; Sangkloy, Burnell, Ham, & Hays, 2016). Second, drawing behavior has been used to characterize the neural mechanisms that support visual understanding and production (Fan, Wammes, et al., 2020), as well as the cognitive consequences of brain injury and disease (Bozeat et al., 2003; Chen & Goedert, 2012; Bozeat, Ralph, Patterson, Garrard, & Hodges, 2000). Third, a combination of large-scale data collection efforts and targeted experimental interventions focused on graphical production have led to more detailed understanding of the processes by which learners learn about the visual world, as well as more abstract concepts (Dillon, 2021; Long et al., 2021; Fiorella & Zhang, 2018).

Our workshop will consist of two complementary components: First, we will host three virtual seminars to introduce each thematic perspective (Artificial Intelligence, Neuroscience, and Developmental Science) *in advance* of the Cognitive Science Society Conference 2022. This 'flipped' format will be more inclusive, by allowing attendees to fully participate from anywhere in the world. Additionally, we will record these seminars so that attendees can also view these talks asynchronously.

Second, we will host an in-person poster session and facilitate group discussion of issues raised during the thematic seminars on site at the conference in July. These activities are inherently active and interactive, and will thus empower in-person attendees to capitalize on this time to interact informally with other researchers (rather than attend a long series of back-to-back talks). We plan to dedicate the funds allocated to this workshop to support early career researchers by awarding two student travel grants and one 'Best Poster Award.' Taken together, we believe our strategy will lay a strong foundation for further interaction among attendees be-

yond CogSci2022, and even inspire new multi-disciplinary collaborations.

Organizers

The organizers share a deep interest in leveraging drawings to gain insight into cognitive phenomena, including perception, semantic memory, natural pedagogy, among others.

Kushin Mukherjee (Organizer) is a PhD student at the University of Wisconsin-Madison. His research focuses on how visual concepts are acquired and reshaped through experience and how people communicate these concepts to each other through visual conduits.

Holly Huey (Organizer) is a PhD student at the University of California, San Diego. Her research investigates how people use visual explanations to convey abstract knowledge to others (e.g., how causal systems work) in order to advance theories of explanatory behavior, as well as explore the role of visualizations in scientific learning and pedagogy.

Timothy T. Rogers (Organizer) is Professor of Psychology at the University of Wisconsin-Madison. Research in his lab investigates human semantic memory — that is, knowledge about the meanings of words, objects, and events.

Judith E. Fan (Organizer) is Assistant Professor of Psychology at the University of California, San Diego. Research in her lab focuses on the use of physical representations of thought, including sketches and other objects, during learning, communication, and problem solving.

Invited Seminar Speakers

For each virtual seminar, speakers will present short talks while the chairs will be responsible for leading the session and providing concluding remarks. We anticipate the first seminar taking place in early June 2022, the second in late June 2022, and the third in early July 2022, depending on speaker and chair availability.

Artificial Intelligence Seminar

David Ha (Chair) is a Research Scientist at Google Brain Tokyo. His work involves reinforcement learning and leveraging deep learning advances to build generative models of sketch production.

Catherine Wong (Speaker) is a PhD student at MIT. Her research uses program synthesis as a framework to understand how people and machines learn conceptual abstractions to flexibly communicate across different contexts.

Yulia Gryaditskaya (Speaker) is a Senior Research Fellow at the Center for Vision, Speech, and Signal Processing at the University of Surrey. Her research uses computer vision techniques, like geometric deep learning, to model and understand sketch generation and classification.

Neuroscience Seminar

Wilma Bainbridge (Chair) is an Assistant Professor of Psychology at the University of Chicago. She studies the neuroscience of perception and memory, and uses behavioral drawing tasks along with functional brain imaging to investigate the visual content of memories.

Jeff Wammes (Speaker) is an Assistant Professor of Psychology at Queen's University. His research focuses on how learning through active learning tasks, such as drawing, can reshape how memories are represented in memories.

Matt Lambon Ralph (Speaker) is the Director of the MRC Cognition and Brain Sciences Unit at the University of Cambridge. His research focuses on the neuropsychology of semantic cognition and aphasia using neuroscientific computational models, TMS, and neuroimaging techniques.

Developmental Science Seminar

Moira Dillon (Chair) is an Assistant Professor of Psychology at New York University. Her work investigates children's drawings to probe their geometric understanding through developmental, cognitive, and computational lenses.

Bria Long (Speaker) is a Postdoctoral Scholar at Stanford University. She studies how children's visual experience of the world scaffolds early learning by investigating large-scale datasets of children's drawings with the help of deep neural networks.

Logan Fiorella (Speaker) is an Associate Professor of Education Psychology at the University of Georgia. His research is on how different kinds of instructional methods, such as the use of graphics and diagrams, can be leveraged to improve learning in STEM.

Acknowledgments

References

Abell, C. (2009). Canny resemblance. *Philosophical Review*, *118*(2), 183–223.

Bainbridge, W. A., Hall, E. H., & Baker, C. I. (2019). Drawings of real-world scenes during free recall reveal detailed object and spatial information in memory. *Nature communications*, *10*(1), 1–13.

Bhunia, A. K., Das, A., Muhammad, U. R., Yang, Y., Hospedales, T. M., Xiang, T., . . . Song, Y.-Z. (2020). Pixelor: A competitive sketching ai agent. so you think you can sketch? *ACM Transactions on Graphics (TOG)*, *39*(6), 1–15.

Bozeat, S., Lambon Ralph, M. A., Graham, K. S., Patterson, K., Wilkin, H., Rowland, J., . . . Hodges, J. R. (2003). A duck with four legs: Investigating the structure of conceptual knowledge using picture drawing in semantic dementia. *Cognitive Neuropsychology*, *20*(1), 27–47.

Bozeat, S., Ralph, M. A. L., Patterson, K., Garrard, P., & Hodges, J. R. (2000). Non-verbal semantic impairment in semantic dementia. *Neuropsychologia*, *38*(9), 1207–1215.

Chamberlain, R. (2018). Drawing as a window onto expertise. *Current Directions in Psychological Science*, *27*(6), 501–507.

Chamberlain, R., & Wagemans, J. (2015). Visual arts training is linked to flexible attention to local and global levels of visual stimuli. *Acta Psychologica*, *161*, 185–197.

Chen, P., & Goedert, K. M. (2012). Clock drawing in spatial neglect: A comprehensive analysis of clock perimeter, placement, and accuracy. *Journal of Neuropsychology*, *6*(2), 270–289.

Clottes, J. (2008). *Cave art*. Phaidon London.

Dillon, M. R. (2021). Rooms without walls: Young children draw objects but not layouts. *Journal of Experimental Psychology: General*, *150*(6), 1071.

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., Wammes, J. D., Gunn, J. B., Yamins, D. L., Norman, K. A., & Turk-Browne, N. B. (2020). Relating visual production and recognition of objects in human visual cortex. *Journal of Neuroscience*, *40*(8), 1710–1721.

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.

Fiorella, L., & Zhang, Q. (2018). Drawing boundary conditions for learning by drawing. *Educational Psychology Review*, *30*(3), 1115–1137.

Gombrich, E. (1989). *The story of art*. Phaidon Press, Ltd.

Goodman, N. (1976). *Languages of art: An approach to a theory of symbols*. Hackett publishing.

Greenberg, G. (2021). Semantics of pictorial space. *Review of Philosophy and Psychology*, 1–41.

Ha, D., & Eck, D. (2017). A neural representation of sketch drawings. *arXiv preprint arXiv:1704.03477*.

Hawkins, R. D., Sano, M., Goodman, N. D., & Fan, J. E. (2021). Visual resemblance and communicative context constrain the emergence of graphical conventions. *arXiv preprint arXiv:2109.13861*.

Hertzmann, A. (2020). Why do line drawings work? a realism hypothesis. *Perception*, *49*(4), 439–451.

Huey, H., Walker, C., & Fan, J. (2021). Explanatory drawings prioritize functional properties at the expense of visual fidelity.

Jongejan, J., Rowley, H., Kawashima, T., Kim, J., & Fox-Gieg, N. (2017). *Google quick, draw*.

Karmiloff-Smith, A. (1990). Constraints on representational change: Evidence from children's drawing. *Cognition*, *34*(1), 57–83.

Kellogg, R. (1969). *Analyzing children's art*. National Press Books Palo Alto, CA.

Long, B., Fan, J., Chai, Z., & Frank, M. C. (2021). Parallel developmental changes in children's drawing and recognition of visual concepts.

Mukherjee, K., Hawkins, R. X., & Fan, J. W. (2019). Communicating semantic part information in drawings. In *Cogsci* (pp. 2413–2419).

Perdreau, F., & Cavanagh, P. (2014). Drawing skill is related to the efficiency of encoding object structure. *i-Perception*, *5*(2), 101–119.

Phillips, W., Hobbs, S., & Pratt, F. (1978). Intellectual realism in children's drawings of cubes. *Cognition*, *6*(1), 15–33.

Sangkloy, P., Burnell, N., Ham, C., & Hays, J. (2016). The sketchy database: learning to retrieve badly drawn bunnies. *ACM Transactions on Graphics (TOG)*, *35*(4), 1–12.

Wammes, J. D., Meade, M. E., & Fernandes, M. A. (2016). The drawing effect: Evidence for reliable and robust memory benefits in free recall. *Quarterly Journal of Experimental Psychology*, *69*(9), 1752–1776.

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

	Workshop Budget		
Item Name	Amount per award	Number of awards	Description
Early Career Researcher Travel Award	$200	2	These awards would support student and/or postdoc presenters to fully participate in the workshop, whether be defraying the cost of travel OR mitigating other barriers to participation (e.g., childcare expenses)
Best Poster Award	$200	1	This award would be presented to the most outstanding poster presentation at the workshop, as determined by the members of an expert panel to be appointed by the workshop organizers.
Total Amount	$600		